

BEST AVAILABLE COPY

JCS30 U.S. PTO
09/477365



Data-Over-Cable Service Interface Specifications

Radio Frequency Interface Specification

SP-RFiv1.1-I02-990731

INTERIM
SPECIFICATION

Notice

This document is a cooperative effort undertaken at the direction of Cable Television Laboratories, Inc. for the benefit of the cable industry in general. Neither CableLabs nor any member company is responsible for any liability of any nature whatsoever resulting from or arising out of use or reliance upon this specification by any party. This document is furnished on an "AS IS" basis and neither CableLabs nor its members provides any representation or warranty, express or implied, regarding its accuracy, completeness, or fitness for a particular purpose. Distribution of this document is restricted pursuant to the terms of separate access agreements negotiated with each of the parties to whom this document has been furnished.

© Copyright 1999 Cable Television Laboratories, Inc.
All rights reserved.

Document Status Sheet

Document Control Number: SP-RFiv1.1-102-990731

Reference: Radio Frequency Interface Specification

Revision History: 101 — First Interim Release, March 11, 1999
102 — Second Interim Release, July 31, 1999

Date: July 31, 1999

Status Code: ~~Work in Process~~ Draft Interim Released

Distribution Restrictions: ~~Cable Lab only~~ ~~CL~~ Reviewers ~~CL~~ Vendor Public

Key to Document Status Codes

- Work In Process** An incomplete document, designed to guide discussion and generate feedback, that may include several alternative requirements for consideration.
- Draft** A document in specification format considered largely complete, but lacking review by cable industry and vendors. Drafts are susceptible to substantial change during the review process.
- Interim** A document which has undergone rigorous cable industry and vendor review, suitable for use by vendors to design in conformance with, and suitable for field testing.
- Released** A stable document, reviewed, tested and validated, suitable to enable cross-vendor interoperability.

Contents

1	SCOPE AND PURPOSE.....	1
1.1	SCOPE.....	1
1.2	REQUIREMENTS.....	1
1.3	BACKGROUND.....	1
1.3.1	Service Goals.....	1
1.3.2	Reference Architecture.....	2
1.3.3	Categories of Interface Specification.....	4
1.3.4	Statement of Compatibility.....	5
2	FUNCTIONAL ASSUMPTIONS.....	7
2.1	BROADBAND ACCESS NETWORK.....	7
2.2	EQUIPMENT ASSUMPTIONS.....	7
2.2.1	Frequency Plan.....	7
2.2.2	Compatibility with Other Services.....	7
2.2.3	Fault Isolation Impact on Other Users.....	8
2.2.4	Cable System Terminal Devices.....	8
2.3	RF CHANNEL ASSUMPTIONS.....	8
2.3.1	Transmission Downstream.....	8
2.3.2	Transmission Upstream.....	9
2.4	TRANSMISSION LEVELS.....	10
2.5	FREQUENCY INVERSION.....	10
3	COMMUNICATION PROTOCOLS.....	11
3.1	PROTOCOL STACK.....	11
3.1.1	CM and CMTS as Hosts.....	11
3.1.2	Data Forwarding Through the CM and CMTS.....	12
3.2	THE MAC FORWARDER.....	15
3.2.1	Rules for Data-Link-Layer Forwarding.....	16
3.3	NETWORK LAYER.....	16
3.3.1	Requirements for IGMP Management.....	17
3.4	ABOVE THE NETWORK LAYER.....	18
3.5	DATA LINK LAYER.....	18
3.5.1	LLC Sublayer.....	19
3.5.2	Link-Layer Security Sublayer.....	19
3.5.3	MAC Sublayer.....	19
3.6	PHYSICAL LAYER.....	19
3.6.1	Downstream Transmission Convergence Sublayer.....	19
3.6.2	PMD Sublayer.....	20
4	PHYSICAL MEDIA DEPENDENT SUBLAYER SPECIFICATION.....	21
4.1	SCOPE.....	21
4.2	UPSTREAM.....	21

4.2.1	Overview.....	21
4.2.2	Modulation Formats.....	22
4.2.3	FEC Encode.....	25
4.2.4	Scrambler (Randomizer).....	25
4.2.5	Preamble Prepend.....	26
4.2.6	Burst Profiles.....	26
4.2.7	Burst Timing Convention.....	29
4.2.8	Transmit Power Requirements.....	30
4.2.9	Fidelity Requirements.....	31
4.2.10	Frame Structure.....	35
4.2.11	Signal Processing Requirements.....	36
4.2.12	Upstream Demodulator Input Power Characteristics.....	37
4.2.13	Upstream Electrical Output from the CM.....	38
4.3	DOWNSTREAM.....	38
4.3.1	Downstream Protocol.....	38
4.3.2	Scalable Interleaving to Support Low Latency.....	38
4.3.3	Downstream Frequency Plan.....	39
4.3.4	CMTS Output Electrical.....	39
4.3.5	Downstream Electrical Input to CM.....	40
4.3.6	CM BER Performance.....	40
4.3.7	CMTS Timestamp Jitter.....	41
5	DOWNSTREAM TRANSMISSION CONVERGENCE SUBLAYER.....	43
5.1	INTRODUCTION.....	43
5.2	MPEG PACKET FORMAT.....	43
5.3	MPEG HEADER FOR DOCSIS DATA-OVER-CABLE.....	44
5.4	MPEG PAYLOAD FOR DOCSIS DATA-OVER-CABLE.....	44
5.5	INTERACTION WITH THE MAC SUBLAYER.....	45
5.6	INTERACTION WITH THE PHYSICAL LAYER.....	46
5.7	MPEG HEADER SYNCHRONIZATION AND RECOVERY.....	46
6	MEDIA ACCESS CONTROL SPECIFICATION.....	47
6.1	INTRODUCTION.....	47
6.1.1	Overview.....	47
6.1.2	Definitions.....	47
6.1.3	Future Use.....	49
6.2	MAC FRAME FORMATS.....	49
6.2.1	Generic MAC Frame Format.....	49
6.2.2	Packet-Based MAC Frames.....	53
6.2.3	ATM Cell MAC Frames.....	55
6.2.4	Reserved PDU MAC Frames.....	55
6.2.5	MAC-Specific Headers.....	56
6.2.6	Extended MAC Headers.....	61

6.2.7	Fragmented MAC Frames.....	65
6.2.8	Error-Handling	67
6.3	MAC MANAGEMENT MESSAGES	68
6.3.1	MAC Management Message Header.....	68
6.3.2	Time Synchronization (SYNC).....	70
6.3.3	Upstream Channel Descriptor (UCD).....	71
6.3.4	Upstream Bandwidth Allocation Map (MAP).....	75
6.3.5	Ranging Request (RNG-REQ).....	78
6.3.6	Ranging Response (RNG-RSP)	79
6.3.7	Registration Request (REG-REQ).....	83
6.3.8	Registration Response (REG-RSP)	85
6.3.9	Registration Acknowledge (REG-ACK).....	88
6.3.10	Upstream Channel Change Request (UCC-REQ)	90
6.3.11	Upstream Channel Change Response (UCC-RSP).....	91
6.3.12	Dynamic Service Addition — Request (DSA-REQ).....	92
6.3.13	Dynamic Service Addition — Response (DSA-RSP).....	94
6.3.14	Dynamic Service Addition — Acknowledge (DSA-ACK).....	96
6.3.15	Dynamic Service Change — Request (DSC-REQ).....	97
6.3.16	Dynamic Service Change — Response (DSC-RSP).....	98
6.3.17	Dynamic Service Change — Acknowledge (DSC-ACK).....	100
6.3.18	Dynamic Service Deletion — Request (DSD-REQ).....	101
6.3.19	Dynamic Service Deletion — Response (DSD-RSP).....	102
7	MEDIA ACCESS CONTROL PROTOCOL OPERATION	103
7.1	UPSTREAM BANDWIDTH ALLOCATION	103
7.1.1	The Allocation Map MAC Management Message.....	104
7.1.2	Information Elements	104
7.1.3	Requests.....	106
7.1.4	Information Element Feature Compatibility Summary.....	107
7.1.5	Map Transmission and Timing.....	107
7.1.6	Protocol Example.....	108
7.2	SUPPORT FOR MULTIPLE CHANNELS	109
7.3	TIMING AND SYNCHRONIZATION	109
7.3.1	Global Timing Reference	110
7.3.2	CM Channel Acquisition	110
7.3.3	Ranging	110
7.3.4	Timing Units and Relationships.....	111
7.4	UPSTREAM TRANSMISSION AND CONTENTION RESOLUTION.....	112
7.4.1	Contention Resolution Overview.....	112
7.4.2	Transmit Opportunities	113
7.4.3	CM Bandwidth Utilization	114
7.5	DATA LINK ENCRYPTION SUPPORT	114
7.5.1	MAC Messages.....	114

7.5.2	<i>Framing</i>	115
8	QUALITY OF SERVICE & FRAGMENTATION	117
8.1	THEORY OF OPERATION	117
8.1.1	<i>Concepts</i>	118
8.1.2	<i>Object Model</i>	122
8.1.3	<i>Service Classes</i>	123
8.1.4	<i>Authorization</i>	124
8.1.5	<i>Types of Service Flows</i>	125
8.1.6	<i>Service Flows and Classifiers</i>	127
8.1.7	<i>General Operation</i>	128
8.2	UPSTREAM SERVICE FLOW SCHEDULING SERVICES	132
8.2.1	<i>Unsolicited Grant Service</i>	132
8.2.2	<i>Real-Time Polling Service</i>	132
8.2.3	<i>Unsolicited Grant Service with Activity Detection</i>	132
8.2.4	<i>Non-Real-Time Polling Service</i>	133
8.2.5	<i>Best Effort Service</i>	133
8.2.6	<i>Other Services</i>	133
8.3	FRAGMENTATION	134
8.3.1	<i>CM Fragmentation Support</i>	134
8.3.2	<i>CMTS Fragmentation Support</i>	136
8.3.3	<i>Fragmentation Example</i>	137
8.4	PAYLOAD HEADER SUPPRESSION.....	140
8.4.1	<i>Overview</i>	140
8.4.2	<i>Example Applications</i>	141
8.4.3	<i>Operation</i>	141
8.4.4	<i>Signaling</i>	143
8.4.5	<i>Payload Header Suppression Examples</i>	145
9	CABLE MODEM - CMTS INTERACTION	147
9.1	CMTS INITIALIZATION	147
9.2	CABLE MODEM INITIALIZATION	147
9.2.1	<i>Scanning and Synchronization to Downstream</i>	149
9.2.2	<i>Obtain Upstream Parameters</i>	150
9.2.3	<i>Message Flows During Scanning and Upstream Parameter Acquisition</i>	152
9.2.4	<i>Ranging and Automatic Adjustments</i>	153
9.2.5	<i>Establish IP Connectivity</i>	157
9.2.6	<i>Establish Time of Day</i>	157
9.2.7	<i>Transfer Operational Parameters</i>	158
9.2.8	<i>Registration</i>	158
9.2.9	<i>Baseline Privacy Initialization</i>	163
9.2.10	<i>Service IDs During CM Initialization</i>	163
9.2.11	<i>Multiple-Channel Support</i>	164

9.3	STANDARD OPERATION	164
9.3.1	Periodic Signal Level Adjustment	164
9.3.2	Changing Upstream Burst Parameters	166
9.3.3	Changing Upstream Channels	167
9.4	DYNAMIC SERVICE	169
9.4.1	Dynamic Service Addition	171
9.4.2	Dynamic Service Change	181
9.4.3	Dynamic Service Deletion	192
9.5	FAULT DETECTION AND RECOVERY	200
9.5.1	Prevention of Unauthorized Transmissions	200
10	SUPPORTING FUTURE NEW CABLE MODEM CAPABILITIES	201
10.1	DOWNLOADING CABLE MODEM OPERATING SOFTWARE	201
	APPENDIX A. WELL-KNOWN ADDRESSES	203
A.1	MAC ADDRESSES	203
A.2	MAC SERVICE IDS	203
A.2.1	All CMs and No CM Service IDs	203
A.2.2	Well-Known 'Multicast' Service IDs	203
A.2.3	Priority Request Service IDs	204
A.3	MPEG PID	204
	APPENDIX B. PARAMETERS AND CONSTANTS	205
	APPENDIX C. COMMON RADIO FREQUENCY INTERFACE ENCODINGS	207
C.1	ENCODINGS FOR CONFIGURATION AND MAC-LAYER MESSAGING	207
C.1.1	Configuration File and Registration Settings	207
C.1.2	Configuration-File-Specific Settings	215
C.1.3	Registration-Request/Response-Specific Encodings	217
C.1.4	Dynamic-Service-Message-Specific Encodings	220
C.2	QUALITY-OF-SERVICE-RELATED ENCODINGS	221
C.2.1	Packet Classification Encodings	221
C.2.2	Service Flow Encodings	229
C.2.3	Parameter Applicability for Upstream Service Scheduling	245
C.3	ENCODINGS FOR OTHER INTERFACES	246
C.3.1	Telephone Settings Option	246
C.3.2	Baseline Privacy Configuration Settings Option	246
C.4	CONFIRMATION CODE	247
	APPENDIX D. CM CONFIGURATION INTERFACE SPECIFICATION	249
D.1	CM IP ADDRESSING	249
D.1.1	DHCP Fields Used by the CM	249
D.2	CM CONFIGURATION	250
D.2.1	CM Binary Configuration File Format	250

D.2.2	Configuration File Settings	251
D.2.3	Configuration File Creation	252
D.3	CONFIGURATION VERIFICATION	253
D.3.1	CMTS MIC Calculation	254
APPENDIX E. MAC SERVICE DEFINITION		257
E.1	MAC SERVICE OVERVIEW	257
E.1.1	MAC Service Parameters	258
E.2	MAC DATA SERVICE INTERFACE	259
E.2.1	MAC_DATA.request	259
E.2.2	MAC_DATA.indicate	260
E.2.3	MAC_GRANT_SYNCHRONIZE.indicate	260
E.2.4	MAC_CMTS_MASTER_CLOCK_SYNCHRONIZE.indicate	261
E.3	MAC CONTROL SERVICE INTERFACE	261
E.3.1	MAC_REGISTRATION_RESPONSE.indicate	261
E.3.2	MAC_CREATE_SERVICE_FLOW.request	261
E.3.3	MAC_CREATE_SERVICE_FLOW.response	262
E.3.4	MAC_CREATE_SERVICE_FLOW.indicate	262
E.3.5	MAC_DELETE_SERVICE_FLOW.request	262
E.3.6	MAC_DELETE_SERVICE_FLOW.response	262
E.3.7	MAC_DELETE_SERVICE_FLOW.indicate	263
E.3.8	MAC_CHANGE_SERVICE_FLOW.request	263
E.3.9	MAC_CHANGE_SERVICE_FLOW.response	263
E.3.10	MAC_CHANGE_SERVICE_FLOW.indicate	263
E.4	MAC SERVICE USAGE SCENARIOS	264
E.4.1	Transmission of PDUs from Upper Layer Service to MAC DATA Service	264
E.4.2	Reception of PDUs to Upper Layer Service from MAC DATA Service	264
E.4.3	Sample Sequence of MAC Control and MAC Data Services	264
APPENDIX F. EXAMPLE PREAMBLE SEQUENCE		267
F.1	INTRODUCTION	267
F.2	EXAMPLE PREAMBLE SEQUENCE	267
APPENDIX G. DOCSIS V1.0/V1.1 INTEROPERABILITY		269
G.1	INTRODUCTION	269
G.2	GENERAL INTEROPERABILITY ISSUES	269
G.2.1	Provisioning	269
G.2.2	Registration	269
G.2.3	Dynamic Service Establishment	270
G.2.4	Fragmentation	270
G.2.5	Multicast Support	270
G.2.6	Upstream Channel Change	270
G.3	HYBRID DEVICES	271
G.4	INTEROPERABILITY & PERFORMANCE	271

APPENDIX H. MULTIPLE UPSTREAM CHANNELS.....	273
APPENDIX I. THE DATA-OVER-CABLE SPANNING TREE PROTOCOL	275
I.1 BACKGROUND	275
I.2 PUBLIC SPANNING TREE	275
I.3 PUBLIC SPANNING TREE PROTOCOL DETAILS	276
I.4 SPANNING TREE PARAMETERS AND DEFAULTS	277
APPENDIX J. ERROR CODES AND MESSAGES	279
APPENDIX K. DOCSIS TRANSMISSION AND CONTENTION RESOLUTION	283
K.1 INTRODUCTION:.....	283
APPENDIX L. IGMP EXAMPLE.....	289
L.1 TRANSITION EVENTS.....	289
APPENDIX M. UNSOLICITED GRANT SERVICES	291
M.1 UNSOLICITED GRANT SERVICE (UGS)	291
M.1.1 Introduction.....	291
M.1.2 Configuration Parameters.....	291
M.1.3 Operation	291
M.1.4 Jitter	292
M.1.5 Synchronization Issues	292
M.2 UNSOLICITED GRANT SERVICE WITH ACTIVITY DETECTION (UGS-AD).....	293
M.2.1 Introduction.....	293
M.2.2 MAC Configuration Parameters	293
M.2.3 Operation	293
M.2.4 Example	294
M.2.5 Talk Spurt Grant Burst.....	295
M.2.6 Admission Considerations.....	296
APPENDIX N. REFERENCES.....	297
APPENDIX O. GLOSSARY.....	301
APPENDIX P. ACKNOWLEDGMENT	313
APPENDIX Q. REVISIONS.....	315
Q.1 ECNS INCLUDED IN SP-RFIV1.1-I02-990731	315

This page intentionally left blank.

Figures

FIGURE 1-1.	TRANSPARENT IP TRAFFIC THROUGH THE DATA-OVER-CABLE SYSTEM	2
FIGURE 1-2.	DATA-OVER-CABLE REFERENCE ARCHITECTURE	3
FIGURE 3-1.	PROTOCOL STACK ON THE RF INTERFACE.....	11
FIGURE 3-2.	DATA FORWARDING THROUGH THE CM AND CMTS.....	12
FIGURE 3-3.	EXAMPLE CONDITION FOR NETWORK LOOPS	13
FIGURE 3-4.	MAC FORWARDER.....	15
FIGURE 4-1.	QPSK SYMBOL MAPPING.....	23
FIGURE 4-2.	16QAM GRAY-CODED SYMBOL MAPPING.....	23
FIGURE 4-3.	16QAM DIFFERENTIAL-CODED SYMBOL MAPPING	23
FIGURE 4-4.	SCRAMBLER STRUCTURE.....	26
FIGURE 4-5.	NOMINAL BURST TIMING	29
FIGURE 4-6.	WORST-CASE BURST TIMING.....	30
FIGURE 4-7.	EXAMPLE FRAME STRUCTURES WITH FLEXIBLE BURST LENGTH MODE.....	35
FIGURE 4-8.	SIGNAL-PROCESSING SEQUENCE	36
FIGURE 4-9.	TDMA UPSTREAM TRANSMISSION PROCESSING	37
FIGURE 5-1.	EXAMPLE OF INTERLEAVING MPEG PACKETS IN DOWNSTREAM	43
FIGURE 5-2.	FORMAT OF AN MPEG PACKET	43
FIGURE 5-3.	PACKET FORMAT WHERE A MAC FRAME IMMEDIATELY FOLLOWS THE POINTER_FIELD.....	45
FIGURE 5-4.	PACKET FORMAT WITH MAC FRAME PRECEDED BY STUFFING BYTES	45
FIGURE 5-5.	PACKET FORMAT SHOWING MULTIPLE MAC FRAMES IN A SINGLE PACKET	45
FIGURE 5-6.	PACKET FORMAT WHERE A MAC FRAME SPANS MULTIPLE PACKETS	46
FIGURE 6-1.	GENERIC MAC FRAME FORMAT.....	49
FIGURE 6-2.	UPSTREAM MAC/PMD CONVERGENCE.....	50
FIGURE 6-3.	MAC HEADER FORMAT	51
FIGURE 6-4.	ETHERNET/802.3 PACKET PDU FORMAT.....	53
FIGURE 6-5.	RESERVED PDU FORMAT.....	55
FIGURE 6-6.	TIMING MAC HEADER.....	56
FIGURE 6-7.	MANAGEMENT MAC HEADER	57
FIGURE 6-8.	REQUEST FRAME FORMAT	58
FIGURE 6-9.	FRAGMENTATION MAC HEADER FORMAT	59
FIGURE 6-10.	CONCATENATION OF MULTIPLE MAC FRAMES.....	60
FIGURE 6-11.	CONCATENATION MAC HEADER FORMAT	60
FIGURE 6-12.	EXTENDED MAC FORMAT	61
FIGURE 6-13.	FRAGMENTATION DETAILS.....	65
FIGURE 6-14.	MAC HEADER AND MAC MANAGEMENT MESSAGE HEADER FIELDS.....	68
FIGURE 6-15.	FORMAT OF PACKET PDU FOLLOWING THE TIMING HEADER.....	70
FIGURE 6-16.	UPSTREAM CHANNEL DESCRIPTOR	71
FIGURE 6-17.	TOP-LEVEL ENCODING FOR A BURST DESCRIPTOR	72
FIGURE 6-18.	EXAMPLE OF UCD ENCODED TLV DATA.....	74
FIGURE 6-19.	MAP FORMAT	75
FIGURE 6-20.	MAP INFORMATION ELEMENT STRUCTURE.....	76

FIGURE 6-21.	PACKET PDU FOLLOWING THE TIMING HEADER.....	78
FIGURE 6-22.	RANGING RESPONSE.....	79
FIGURE 6-23.	GENERALIZED DECISION FEEDBACK EQUALIZATION COEFFICIENTS.....	81
FIGURE 6-24.	CMTS DEMODULATOR EQUALIZER TAP LOCATION DEFINITION.....	81
FIGURE 6-25.	EXAMPLE OF TLV DATA.....	82
FIGURE 6-26.	REGISTRATION REQUEST.....	83
FIGURE 6-27.	REGISTRATION RESPONSE FORMAT.....	85
FIGURE 6-28.	REGISTRATION ACKNOWLEDGMENT.....	88
FIGURE 6-29.	UPSTREAM CHANNEL CHANGE REQUEST.....	90
FIGURE 6-30.	UPSTREAM CHANNEL CHANGE RESPONSE.....	91
FIGURE 6-31.	DYNAMIC SERVICE ADDITION — REQUEST.....	92
FIGURE 6-32.	DYNAMIC SERVICE ADDITION — RESPONSE.....	94
FIGURE 6-33.	DYNAMIC SERVICE ADDITION — ACKNOWLEDGE.....	96
FIGURE 6-34.	DYNAMIC SERVICE CHANGE — REQUEST.....	97
FIGURE 6-35.	DYNAMIC SERVICE CHANGE — RESPONSE.....	98
FIGURE 6-36.	DYNAMIC SERVICE CHANGE — ACKNOWLEDGE.....	100
FIGURE 6-37.	DYNAMIC SERVICE DELETION — REQUEST.....	101
FIGURE 6-38.	DYNAMIC SERVICE DELETION — RESPONSE.....	102
FIGURE 7-1.	ALLOCATION MAP.....	103
FIGURE 7-2.	PROTOCOL EXAMPLE.....	108
FIGURE 8-1.	PROVISIONED AUTHORIZATION MODEL “ENVELOPES”.....	119
FIGURE 8-2.	DYNAMIC AUTHORIZATION MODEL “ENVELOPES”.....	120
FIGURE 8-3.	CLASSIFICATION WITHIN THE MAC LAYER.....	121
FIGURE 8-4.	THEORY OF OPERATION OBJECT MODEL.....	123
FIGURE 8-5.	REGISTRATION MESSAGE FLOW.....	128
FIGURE 8-6.	DYNAMIC SERVICE ADDITION MESSAGE FLOW — CM INITIATED.....	130
FIGURE 8-7.	DYNAMIC SERVICE ADDITION MESSAGE FLOW — CMTS INITIATED.....	131
FIGURE 8-8.	CM FRAGMENTATION FLOWCHART.....	135
FIGURE 8-9.	EXAMPLE OF FRAGMENTING A SINGLE PACKET.....	138
FIGURE 8-10.	FRAGMENTED CONCATENATED PACKET EXAMPLE.....	139
FIGURE 8-11.	PAYLOAD HEADER SUPPRESSION OPERATION.....	142
FIGURE 8-12.	PAYLOAD HEADER SUPPRESSION WITH MASKING.....	143
FIGURE 8-13.	PAYLOAD HEADER SUPPRESSION SIGNALING EXAMPLE.....	144
FIGURE 8-14.	UPSTREAM PAYLOAD HEADER SUPPRESSION EXAMPLE.....	145
FIGURE 8-15.	DOWNSTREAM PAYLOAD HEADER SUPPRESSION EXAMPLE.....	146
FIGURE 9-1.	CM INITIALIZATION OVERVIEW.....	148
FIGURE 9-2.	SDL NOTATION.....	149
FIGURE 9-3.	OBTAINING UPSTREAM PARAMETERS.....	151
FIGURE 9-4.	MESSAGE FLOWS DURING SCANNING AND UPSTREAM PARAMETER ACQUISITION.....	152
FIGURE 9-5.	RANGING AND AUTOMATIC ADJUSTMENTS PROCEDURE.....	153
FIGURE 9-6.	INITIAL RANGING - CM.....	154
FIGURE 9-7.	INITIAL RANGING - CM (CONTINUED).....	155
FIGURE 9-8.	INITIAL RANGING - CMTS (FIG. EDITED PER RFI-N-99054 06/29/99. EW).....	156

FIGURE 9-9.	ESTABLISHING IP CONNECTIVITY	157
FIGURE 9-10.	ESTABLISHING TIME OF DAY	158
FIGURE 9-11.	REGISTRATION — CM.....	159
FIGURE 9-12.	WAIT FOR REGISTRATION RESPONSE — CM.....	160
FIGURE 9-13.	REGISTRATION — CMTS (FIGURE EDITED PER RFI-N-99054 06/30/99.EW).....	162
FIGURE 9-14.	REGISTRATION ACKNOWLEDGMENT— CMTS	163
FIGURE 9-15.	PERIODIC RANGING - CMTS	165
FIGURE 9-16.	PERIODIC RANGING - CM VIEW.....	166
FIGURE 9-17.	CHANGING UPSTREAM CHANNELS: CMTS VIEW	167
FIGURE 9-18.	CHANGING UPSTREAM CHANNELS: CM VIEW.....	168
FIGURE 9-19.	DYNAMIC SERVICE FLOW OVERVIEW	169
FIGURE 9-20.	DYNAMIC SERVICE FLOW STATE TRANSITION DIAGRAM.....	170
FIGURE 9-21.	DYNAMIC SERVICE ADDITION INITIATED FROM CM (FIGURE EDITED PER RFI-N-99048 06/30/99. EW)171	
FIGURE 9-22.	DYNAMIC SERVICE ADDITION INITIATED FROM CMTS (FIGURE EDITED PER RFI-N-99048 06/30/ 99. EW)172	
FIGURE 9-23.	CM START STATE (DSA TRANSACTIONS) (FIG REPLACED 06/22/99 RFI-N-99043.EW)	173
FIGURE 9-24.	CM DSA-RSP PENDING STATE (DSA TRANSACTIONS).....	174
FIGURE 9-25.	CM DSA-ACK PENDING STATE (DSA TRANSACTIONS)	175
FIGURE 9-26.	CM SF_OPERATIONAL STATE (DSA TRANSACTIONS)	176
FIGURE 9-27.	CMTS START STATE (DSA TRANSACTIONS)	177
FIGURE 9-28.	CMTS DSA-RSP PENDING STATE (DSA TRANSACTIONS).....	178
FIGURE 9-29.	CMTS DSA-ACK PENDING STATE (DSA TRANSACTIONS).....	179
FIGURE 9-30.	CMTS SF_OPERATIONAL STATE (DSA TRANSACTIONS)	180
FIGURE 9-31.	CM-INITIATED DSC	182
FIGURE 9-32.	CMTS-INITIATED DSC	183
FIGURE 9-33.	CM SF_OPERATIONAL STATE (DSC TRANSACTIONS)	184
FIGURE 9-34.	CMDSC-RSPENDINGSTATE(DSCTRANSACTIONS)(FIGUREEDITED06/22/99PERRFI-N-99056.EW) 185	
FIGURE 9-35.	CM DSC-ACK PENDING STATE (DSC TRANSACTIONS).....	186
FIGURE 9-36.	CM START STATE (DSC TRANSACTIONS)	187
FIGURE 9-37.	CMTS SF_OPERATIONAL STATE (DSC TRANSACTIONS).....	188
FIGURE 9-38.	CMTS DSC-RSP PENDING STATE (DSC TRANSACTIONS) (FIGURE EDITED 06/22/99 PER RFI-N- 99056. EW)189	
FIGURE 9-39.	CMTS DSC-ACK PENDING STATE (DSC TRANSACTIONS)	190
FIGURE 9-40.	CMTS START STATE (DSC TRANSACTIONS)	191
FIGURE 9-41.	DYNAMIC SERVICE DELETION INITIATED FROM CM	192
FIGURE 9-42.	DYNAMIC SERVICE DELETION INITIATED FROM CMTS	193
FIGURE 9-43.	CM SF_OPERATIONAL STATE (DSD TRANSACTIONS)	194
FIGURE 9-44.	CMDSD-RSPENDINGSTATE(DSDTRANSACTIONS)(FIGUREEDITED06/22/99RFI-N-99043EW) 195	
FIGURE 9-45.	CM START STATE (DSD TRANSACTIONS) (TITLE EDITED 06/22/99 RFI-N-99043 EW).....	196
FIGURE 9-46.	CMTS SF_OPERATIONAL STATE (DSD TRANSACTIONS)	197
FIGURE 9-47.	CMTS DSD-RSP PENDING STATE (DSD TRANSACTIONS) (FIGURE EDITED 06/22/99 RFI-N-	

99043 EW)198

FIGURE 9-48. CMTS START STATE (DSD TRANSACTIONS) 199

This page intentionally left blank.

Tables

TABLE 2-1.	ASSUMED DOWNSTREAM RF CHANNEL TRANSMISSION CHARACTERISTICS.....	9
TABLE 2-2.	ASSUMED UPSTREAM RF CHANNEL TRANSMISSION CHARACTERISTICS.....	10
TABLE 4-1.	I/Q MAPPING	22
TABLE 4-2.	DERIVATION OF CURRENTLY TRANSMITTED SYMBOL QUADRANT	24
TABLE 4-3.	MAXIMUM CHANNEL WIDTH	24
TABLE 4-4.	BURST PROFILE ATTRIBUTES	27
TABLE 4-5.	USER UNIQUE BURST PARAMETERS.....	27
TABLE 4-6.	SPURIOUS EMISSIONS	32
TABLE 4-7.	ADJACENT CHANNEL SPURIOUS EMISSIONS RELATIVE TO THE TRANSMITTED BURST POWER LEVEL 33	33
TABLE 4-8.	SPURIOUS EMISSIONS IN 5 TO 42 MHz RELATIVE TO THE TRANSMITTED BURST POWER LEVEL 33	33
TABLE 4-9.	MAXIMUM RANGE OF COMMANDED NOMINAL RECEIVE POWER IN EACH CARRIER37	37
TABLE 4-10.	ELECTRICAL OUTPUT FROM CM	38
TABLE 4-11.	INTERLEAVER CHARACTERISTICS	38
TABLE 4-12.	CMTS OUTPUT.....	39
TABLE 4-13.	ELECTRICAL INPUT TO CM	40
TABLE 5-1.	MPEG HEADER FORMAT FOR DOCSIS DATA-OVER-CABLE PACKETS.....	44
TABLE 6-1.	GENERIC MAC HEADER FORMAT	51
TABLE 6-2.	FC FIELD FORMAT	52
TABLE 6-3.	EXAMPLE PACKET PDU FORMAT	54
TABLE 6-4.	EXAMPLE RESERVED PDU FORMAT	55
TABLE 6-5.	MAC-SPECIFIC HEADERS AND FRAMES.....	56
TABLE 6-6.	TIMING MAC HEADER FORMAT	57
TABLE 6-7.	EXAMPLE MANAGEMENT MAC HEADER FORMAT.....	57
TABLE 6-8.	REQUEST FRAME (REQ) FORMAT	58
TABLE 6-9.	FRAGMENTATION MAC FRAME (FRAG) FORMAT	59
TABLE 6-10.	CONCATENATED MAC FRAME FORMAT	60
TABLE 6-11.	EXAMPLE EXTENDED HEADER FORMAT	61
TABLE 6-12.	EH ELEMENT FORMAT	62
TABLE 6-13.	EXTENDED HEADER TYPES	62
TABLE 6-14.	FRAGMENTATION EXTENDED HEADER FORMAT	63
TABLE 6-15.	PAYLOAD HEADER SUPPRESSION EHDR SUB-ELEMENT FORMAT.....	64
TABLE 6-16.	UNSOLICITED GRANT SYNCHRONIZATION EHDR SUB-ELEMENT FORMAT	65
TABLE 6-17.	MAC MANAGEMENT MESSAGE TYPES	69
TABLE 6-18.	CHANNEL TLV PARAMETERS	72
TABLE 6-19.	UPSTREAM PHYSICAL-LAYER BURST ATTRIBUTES	73
TABLE 6-20.	ALLOCATION MAP INFORMATION ELEMENTS (IE)	77
TABLE 6-21.	RANGING RESPONSE MESSAGE ENCODINGS	80
TABLE 7-1.	IE FEATURE COMPATIBILITY SUMMARY.....	107
TABLE 7-2.	EXAMPLE RELATING MINI-SLOTS TO TIME TICKS.....	112

TABLE 7-3.	TRANSMIT OPPORTUNITY	114
TABLE 8-1.	TFTP FILE CONTENTS	129
TABLE 8-2.	REGISTRATION REQUEST CONTENTS	129
TABLE 8-3.	REGISTRATION RESPONSE CONTENTS	130
TABLE 8-4.	PAYLOAD HEADER SUPPRESSION DEFINITIONS	140
TABLE 9-1.	RECOVERY PROCESS ON LOSS OF SPECIFIC MAC MESSAGES	200

Radio Frequency Interface Specifications

1 Scope and Purpose

1.1 Scope

This Interim Specification defines the radio-frequency interface specifications for high-speed data-over-cable systems that are being developed by Cable Television Laboratories for the benefit of the cable industry. It is being issued to facilitate design and field testing leading to the early manufacturability and interoperability of conforming hardware by multiple vendors.

1.2 Requirements

Throughout this document, the words that are used to define the significance of particular requirements are capitalized. These words are:

"MUST"	This word or the adjective "REQUIRED" means that the item is an absolute requirement of this specification.
"MUST NOT"	This phrase means that the item is an absolute prohibition of this specification.
"SHOULD"	This word or the adjective "RECOMMENDED" means that there may exist valid reasons in particular circumstances to ignore this item, but the full implications should be understood and the case carefully weighed before choosing a different course.
"SHOULD NOT"	This phrase means that there may exist valid reasons in particular circumstances when the listed behavior is acceptable or even useful, but the full implications should be understood and the case carefully weighed before implementing any behavior described with this label.
"MAY"	This word or the adjective "OPTIONAL" means that this item is truly optional. One vendor may choose to include the item because a particular marketplace requires it or because it enhances the product, for example; another vendor may omit the same item.

Other text is descriptive or explanatory.

1.3 Background

1.3.1 Service Goals

Cable operators are interested in deploying high-speed packet-based communications systems on cable television systems that are capable of supporting a wide variety of services. Services under consideration by cable operators include packet telephony service, video conferencing service, T1/frame relay equivalent service, and many others. To this end, CableLabs' member companies have decided to prepare a series of interface specifications that will permit the early definition, design, development and deployment of data-over-cable systems on a uniform, consistent, open, non-proprietary, multi-vendor interoperable basis.

The intended service will allow transparent bi-directional transfer of Internet Protocol (IP) traffic, between the cable system headend and customer locations, over an all-coaxial or hybrid-fiber/coax (HFC) cable network. This is shown in simplified form in Figure 1-1.

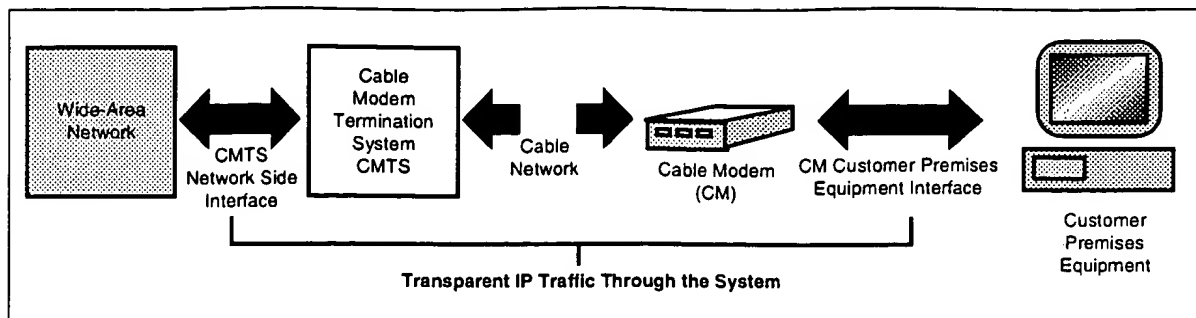


Figure 1-1. Transparent IP Traffic Through the Data-Over-Cable System

The transmission path over the cable system is realized at the headend by a Cable Modem Termination System (CMTS), and at each customer location by a Cable Modem (CM). At the headend (or hub), the interface to the data-over-cable system is called the Cable Modem Termination System - Network-Side Interface (CMTS-NSI) and is specified in [DOCSIS3]. At the customer locations, the interface is called the cable-modem-to-customer-premises-equipment interface (CMCI) and is specified in [DOCSIS4]. The intent is for operators to transparently transfer IP traffic between these interfaces, including but not limited to datagrams, DHCP, ICMP, and IP Group addressing (broadcast and multicast).

1.3.2 Reference Architecture

The reference architecture for the data-over-cable services and interfaces is shown in Figure 1-2.




Figure 1-2. Data-Over-Cable Reference Architecture

1.3.3 Categories of Interface Specification

The basic reference architecture of Figure 1-2 involves three categories of interface.

Data Interfaces - These are the CMCI [DOCSIS4] and CMTS-NSI [DOCSIS3], corresponding respectively to the cable-modem-to-customer-premises-equipment (CPE) interface (for example, between the customer's computer and the cable modem), and the cable modem termination system network-side interface between the cable modem termination system and the data network.

Operations Support Systems Interfaces - These are network element management layer interfaces between the network elements and the high-level OSSs (operations support systems) which support the basic business processes, and are documented in [DOCSIS5].

Telephone Return Interface - CMTRI - This is the interface between the cable modem and a telephone return path, for use in cases where the return path is not provided or not available via the cable network, and is documented in [DOCSIS6].

RF Interfaces - The RF interfaces defined in this document are the following:

- Between the cable modem and the cable network.
- Between the CMTS and the cable network, in the downstream direction (traffic toward the customer)
- Between the CMTS and the cable network, in the upstream direction (traffic from the customer).

Security Requirements -

- Baseline data-over-cable security is defined in [DOCSIS8].

1.3.3.1 Data-Over-Cable Service Interface Documents

A list of the documents in the Data-Over-Cable Service Interface Specifications family is provided below. For updates, please refer to URL <http://www.cablemodem.com>.

Designation	Title
SP-CMCI	Cable Modem to Customer Premises Equipment Interface Specification
SP-CMTS-NSI	Cable Modem Termination System Network Side Interface Specification
SP-CMTRI	Cable Modem Telco Return Interface Specification
SP-OSSI	Operations Support System Interface Specification
SP-RFI	Radio Frequency Interface Specification
SP-BPI+	Baseline Privacy Plus Interface Specification

Key to Designations:

SP	Specification
TP	Test Plan — a document of test procedures to validate specification conformance, interoperability or performance
TR	Technical Report (provides a context for understanding and applying the specification or initial ideas about possible future features)

1.3.4 Statement of Compatibility

This document specifies an interface, commonly referred to as DOCSIS 1.1, which is an extension of the interface specified in [DOCSIS9], commonly referred to as DOCSIS 1.0. These extensions are entirely backwards and forwards compatible with the previous specification. DOCSIS 1.1 compliant CMs **MUST** interoperate seamlessly with DOCSIS 1.0 CMTSs. DOCSIS 1.1 compliant CMTSs **MUST** seamlessly support DOCSIS 1.0 CMs.

Refer to Appendix G for further interoperability information.

This page intentionally left blank.

2 Functional Assumptions

This section describes the characteristics of cable television plant to be assumed for the purpose of operating a data-over-cable system. It is not a description of CMTS or CM parameters. The data-over-cable system **MUST** be interoperable within the environment described in this section.

2.1 Broadband Access Network

A coaxial-based broadband access network is assumed. This may take the form of either an all-coax or hybrid-fiber/coax (HFC) network. The generic term "cable network" is used here to cover all cases.

A cable network uses a shared-medium, tree-and-branch architecture with analog transmission. The key functional characteristics assumed in this document are the following:

- Two-way transmission.
- A maximum optical/electrical spacing between the CMTS and the most distant CM of 100 miles, although typical maximum separation may be 10-15 miles.
- A maximum differential optical/electrical spacing between the CMTS and the closest and most distant modems of 100 miles, although this would typically be limited to 15 miles.

2.2 Equipment Assumptions

2.2.1 Frequency Plan

In the downstream direction, the cable system is assumed to have a passband with a lower edge between 50 and 54 MHz and an upper edge that is implementation-dependent but is typically in the range of 300 to 864 MHz. Within that passband, NTSC analog television signals in 6-MHz channels are assumed to be present on the standard, HRC or IRC frequency plans of [EIA-S542], as well as other narrowband and wideband digital signals.

In the upstream direction, the cable system may have a subsplit (5-30 MHz) or extended subsplit (5-40 or 5-42 MHz) passband. NTSC analog television signals in 6-MHz channels may be present, as well as other signals.

2.2.2 Compatibility with Other Services

The CM and CMTS **MUST** coexist with the other services on the cable network. In particular,

- a) They **MUST** be interoperable in the cable spectrum assigned for CMTS-CM interoperation while the balance of the cable spectrum is occupied by any combination of television and other signals; and
- b) They **MUST NOT** cause harmful interference to any other services that are assigned to the cable network in spectrum outside of that allocated to the CMTS.

The latter is understood as

- No measurable degradation (highest level of compatibility),
- No degradation below the perceptible level of impairments for all services (standard or medium level of compatibility), or
- No degradation below the minimal standards accepted by the industry (for example, FCC for analog video services) or other service provider (minimal level of compatibility).

2.2.3 Fault Isolation Impact on Other Users

As the data-over-cable system is a shared-media, point-to-multipoint system, fault-isolation procedures should take into account the potential harmful impact of faults and fault-isolation procedures on numerous users of the data-over-cable and other services.

For the interpretation of harmful impact, see Section 2.2.2 above.

2.2.4 Cable System Terminal Devices

The CM **MUST** meet and **SHOULD** exceed all applicable regulations for Cable System Termination Devices and Cable Ready Consumer Equipment as defined in FCC Part 15 [FCC15] and Part 76 [FCC76]. None of these specific requirements may be used to relax any of the specifications contained elsewhere within this document.

2.3 RF Channel Assumptions

The data-over-cable system, configured with at least one set of defined physical-layer parameters (e.g., modulation, forward error correction, symbol rate, etc.) from the range of configuration settings described in this specification, **MUST** be interoperable on cable networks having characteristics defined in this section in such a manner that the forward error correction provides for equivalent operation in a cable system both with and without the impaired channel characteristics described below.

2.3.1 Transmission Downstream

The RF channel transmission characteristics of the cable network in the downstream direction are described in Table 2-1. These numbers assume total average power of a digital signal in a 6-MHz channel bandwidth for carrier levels unless indicated otherwise. For impairment levels, the numbers in Table 2-1 assume average power in a bandwidth in which the impairment levels are measured in a standard manner for cable TV system. For analog signal levels, the numbers in Table 2-1 assume peak envelope power in a 6-MHz channel bandwidth. All conditions are present concurrently. No combination of the following parameters will exceed any stated interface limit defined elsewhere in this specification.

Table 2-1. Assumed Downstream RF Channel Transmission Characteristics

Parameter	Value
Frequency range	Cable system normal downstream operating range is from 50 MHz to as high as 860 MHz. However, the values in this table apply only at frequencies ≥ 88 MHz.
RF channel spacing (design bandwidth)	6 MHz
Transit delay from headend to most distant customer	≤ 0.800 msec (typically much less)
Carrier-to-noise ratio in a 6-MHz band	Not less than 35 dB ³
Carrier-to-Composite triple beat distortion ratio	Not less than 41 dB ³
Carrier-to-Composite second order distortion ratio	Not less than 41 dB ³
Carrier-to-Cross-modulation ratio	Not less than 41 dB ³
Carrier-to-any other discrete interference (ingress)	Not less than 41 dB ³
Amplitude ripple	3 dB within the design bandwidth
Group delay ripple in the spectrum occupied by the CMTS	75 ns within the design bandwidth
Micro-reflections bound for dominant echo	-20 dBc @ ≤ 1.5 μ sec, -30 dBc @ > 1.5 μ sec -10 dBc @ ≤ 0.5 μ sec, -15 dBc @ ≤ 1.0 μ sec
Carrier hum modulation	Not greater than -26 dBc (5%)
Burst noise	Not longer than 25 μ sec at a 10 Hz average rate
Maximum analog video carrier level at the CM input	17 dBmV
Maximum number of analog carriers	121

Notes to Table 2-1:

1. Transmission is from the headend combiner to the CM input at the customer location.
2. Measurement methods defined in [NCTA] or [CableLabs2].
3. Measured relative to a QAM signal that is equal to the nominal video level in the plant.

2.3.2 Transmission Upstream

The RF channel transmission characteristics of the cable network in the upstream direction are described in Table 2-2. All conditions are present concurrently. No combination of the following parameters will exceed any stated interface limit defined elsewhere in this specification.

Table 2-2. Assumed Upstream RF Channel Transmission Characteristics

Parameter	Value
Frequency range	5 to 42 MHz edge to edge
Transit delay from the most distant CM to the nearest CM or CMTS	≤ 0.800 msec (typically much less)
Carrier-to-interference plus ingress (the sum of noise, distortion, common-path distortion and cross-modulation and the sum of discrete and broadband ingress signals, impulse noise excluded) ratio	Not less than 25 dB (Note 2)
Carrier hum modulation	Not greater than -23 dBc (7.0%)
Burst noise	Not longer than 10 μ sec at a 1 kHz average rate for most cases (Notes 3 and 4)
Amplitude ripple 5-42 MHz:	0.5 dB/MHz
Group delay ripple 5-42 MHz:	200 ns/MHz
Micro-reflections -- single echo	-10 dBc @ ≤ 0.5 μ sec -20 dBc @ ≤ 1.0 μ sec -30 dBc @ > 1.0 μ sec
Seasonal and diurnal reverse gain (loss) variation	Not greater than 14 dB min to max

Notes to Table 2-1:

1. Transmission is from the CM output at the customer location to the headend.
2. Ingress avoidance or tolerance techniques MAY be used to ensure operation in the presence of time-varying discrete ingress signals that could be as high as 10 dBc. The ratios are guaranteed only within the digital carrier channels.
3. Amplitude and frequency characteristics sufficiently strong to partially or wholly mask the data carrier.
4. Impulse noise levels more prevalent at lower frequencies (< 15 MHz).

2.3.2.1 Availability

Typical cable network availability is considerably greater than 99%.

2.4 Transmission Levels

The nominal power level of the downstream CMTS signal(s) within a 6-MHz channel is targeted to be in the range -10 dBc to -6 dBc relative to analog video carrier level and will normally not exceed analog video carrier level. The nominal power level of the upstream CM signal(s) will be as low as possible to achieve the required margin above noise and interference. Uniform power loading per unit bandwidth is commonly followed in setting upstream signal levels, with specific levels established by the cable network operator to achieve the required carrier-to-noise and carrier-to-interference ratios.

2.5 Frequency Inversion

There will be no frequency inversion in the transmission path in either the downstream or upstream directions, i.e., a positive change in frequency at the input to the cable network will result in a positive change in frequency at the output.

3 Communication Protocols

This section provides a high-level overview of the communication protocols that **MUST** be used in the data-over-cable system. Detailed specifications for the physical media dependent, downstream transmission, and media access control sublayers are provided in Section 4, Section 5, and Section 6, respectively.

3.1 Protocol Stack

The CM and CMTS operate as forwarding agents and also as end-systems (hosts). The protocol stacks used in these modes differ as shown below.

The principal function of the cable modem system is to transmit Internet Protocol (IP) packets transparently between the headend and the subscriber location. Certain management functions also ride on IP, so that the protocol stack on the cable network is as shown in Figure 3-1 (this does not restrict the generality of IP transparency between the headend and the customer). These management functions include, for example, supporting spectrum management functions and the downloading of software.

3.1.1 CM and CMTS as Hosts

CMs and CMTSs will operate as IP and LLC hosts in terms of IEEE Standard 802 [IEEE802] for communication over the cable network. The protocol stack at the CM and CMTS RF interfaces is shown in Figure 3-1.




Figure 3-1. Protocol Stack on the RF Interface

The CM and CMTS **MUST** function as IP hosts. As such, the CM and CMTS **MUST** support IP and ARP over DIX link-layer framing (see [DIX]). The CMTS **MUST NOT** transmit frames that are smaller than the DIX 64 byte minimum on a downstream channel.¹ However, the CM **MAY** transmit frames that are smaller than the DIX 64 byte minimum on an upstream channel.

The CM and CMTS **MAY** also support IP and ARP over SNAP framing [RFC-1042].

The CM and CMTS also **MUST** function as LLC hosts. As such, the CM and CMTS **MUST** respond appropriately to TEST and XID requests per [ISO8802-2].

3.1.2 Data Forwarding Through the CM and CMTS

3.1.2.1 General

Data forwarding through the CMTS **MAY** be transparent bridging², or **MAY** employ network-layer forwarding (routing, IP switching) as shown in Figure 3-2.

Data forwarding through the CM is link-layer transparent bridging, as shown in Figure 3-2. Forwarding rules are similar to [ISO/IEC10038] with the modifications described in Section 3.1.2.2 and Section 3.1.2.3. This allows the support of multiple network layers.

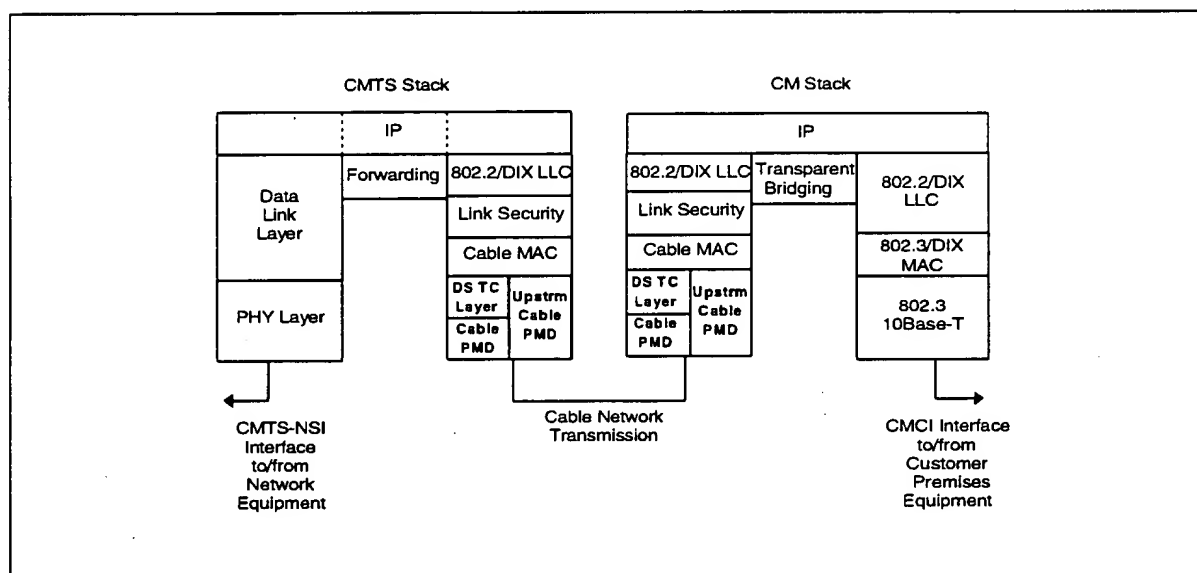


Figure 3-2. Data Forwarding Through the CM and CMTS

Forwarding of IP traffic **MUST** be supported. Support of other network layer protocols is **OPTIONAL**. The ability to restrict the network layer to a single protocol such as IP is **REQUIRED**.

Support for the 802.1d spanning tree protocol of [ISO/IEC10038] with the modifications described in Appendix I is **OPTIONAL** for CMs intended for residential use. CMs intended for commercial use and bridging CMTSs

1. Except as a result of Payload Header Suppression. Refer to Section 8.4.

2. With the exception that for packet PDUs less than 64 bytes to be forwarded from the upstream RFI, a CMTS **MUST** pad out the packet PDU and recompute the CRC.

MUST support this version of spanning tree. CMs and CMTSs MUST include the ability to filter (and disregard) 802.1d BPDUs.

This specification assumes that CMs intended for residential use will not be connected in a configuration which would create network loops such as that shown in Figure 3-3.

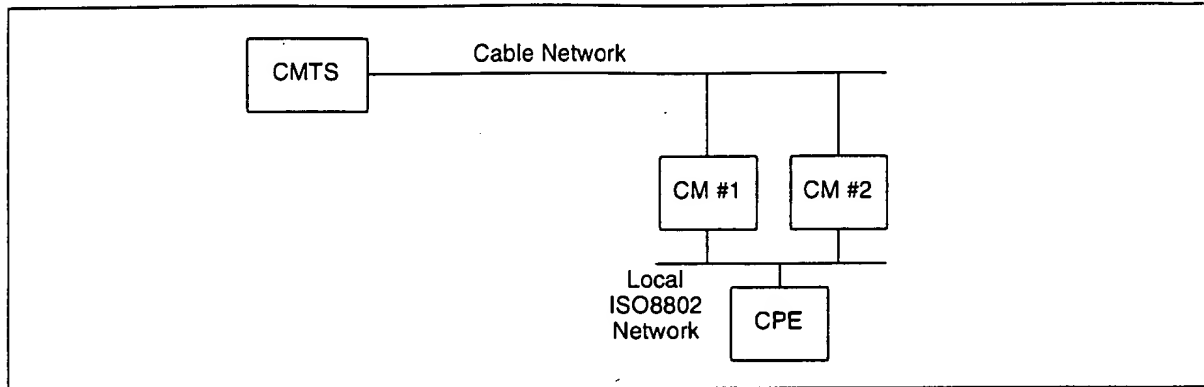


Figure 3-3. Example Condition for Network Loops

3.1.2.2 CMTS Forwarding Rules

At the CMTS, if link-layer forwarding is used, then it MUST conform to the following general 802.1d guidelines:

- Link-layer frames MUST NOT be duplicated.
- Stale frames (those that cannot be delivered in a timely fashion) MUST be discarded.
- Link-layer frames, on a given Service Flow (refer to Section 6.1.2.3), MUST be delivered in the order they are received.

The address-learning and -aging mechanisms used are vendor-dependent.

If network-layer forwarding is used, then the CMTS should conform to IETF Router Requirements [RFC-1812] with respect to its CMTS-RFI and CMTS-NSI interfaces.

Conceptually, the CMTS forwards data packets at two abstract interfaces: between the CMTS-RFI and the CMTS-NSI, and between the upstream and downstream channels. The CMTS MAY use any combination of link-layer (bridging) and network-layer (routing) semantics at each of these interfaces. The methods used at the two interfaces need not be the same.

Forwarding between the upstream and downstream channels within a MAC layer differs from traditional LAN forwarding in that:

- A single channel is simplex, and cannot be considered a complete interface for most protocol (e.g., 802.1d spanning tree, Routing Information Protocol per [RFC-1058]) purposes.
- Upstream channels are essentially point-to-point, whereas downstream channels are shared-media.
- Policy decisions may override full connectivity.

For these reasons, an abstract entity called the MAC Forwarder exists within the CMTS to provide connectivity between stations within a MAC domain (see Section 3.2).

3.1.2.3 CM Forwarding Rules

Data forwarding through the CM is link-layer bridging with the following specific rules.

3.1.2.3.1 CPE MAC Address Acquisition

- The CM **MUST** acquire Ethernet MAC addresses of connected CPE devices, either from the provisioning process or from learning, until the CM acquires its maximum number of CPE MAC addresses (a device-dependent value). Once the CM acquires its maximum number of CPE MAC addresses, then newly discovered CPE MAC addresses **MUST NOT** replace previously acquired addresses. The CM must support acquisition of at least one CPE MAC address.
- The CM **MUST** allow configuration of CPE addresses during the provisioning process (up to its maximum number of CPE addresses) to support configurations in which learning is not practical nor desired.
- Addresses provided during the CM provisioning **MUST** take precedence over learned addresses.
- CPE addresses **MUST NOT** be aged out.
- In order to allow modification of user MAC addresses or movement of the CM, addresses are not retained in non-volatile storage. On a CM reset (e.g. power cycle), all provisioned and learned addresses **MUST** be discarded.

3.1.2.3.2 Forwarding

CM forwarding in both directions **MUST** conform to the following general 802.1d guidelines:

- Link-layer frames **MUST NOT** be duplicated.
- Stale frames (those that cannot be delivered in a timely fashion) **MUST** be discarded.
- Link-layer frames, on a given Service Flow (refer to Section 6.1.2.3), **MUST** be delivered in the order they are received.

Cable-Network-to-Ethernet forwarding **MUST** follow the following specific rules:

- Frames addressed to unknown destinations **MUST NOT** be forwarded from the cable port to the Ethernet port.
- Broadcast frames **MUST** be forwarded to the Ethernet port, unless they are from source addresses which are provisioned or learned as supported CPE devices, in which case they **MUST NOT** be forwarded.
- The forwarding of multicast is controlled by administratively set parameters for the policy-filter service and by a specific multicast tracking algorithm (refer to Section 3.3.1). Multicast frames **MUST NOT** be forwarded unless both mechanisms are in a permissive state.

Ethernet-to-Cable-Network forwarding **MUST** follow the following specific rules:

- Frames addressed to unknown destinations **MUST** be forwarded from the Ethernet port to the cable port.
- Broadcast frames **MUST** be forwarded to the cable port.
- Multicast frames **MUST** be forwarded to the cable port in accordance with filtering configuration settings specified by the cable operator's operations and business support systems.
- Frames from source addresses other than those provisioned or learned as supported CPE devices **MUST NOT** be forwarded.
- If a single-user CM has acquired a MAC address (see Section 3.1.2.3.1), it **MUST NOT** forward data from a second source. Other (non-supported) CPE source addresses **MUST** be learned from the Ethernet port and this information used to filter local traffic as in a traditional learning bridge.

- If a single-user CM has acquired MAC address A as its supported CPE device and learned B as a second device connected to the Ethernet port, it **MUST** filter any traffic from A to B.

3.2 The MAC Forwarder

The MAC Forwarder is a MAC sublayer that resides on the CMTS just below the MAC service access point (MSAP) interface, as shown in Figure 3-4. It is responsible for delivering upstream frames to

- One or more downstream channels
- The MSAP interface.

In Figure 3-4, the LLC sublayer and link security sublayers of the upstream and downstream channels on the cable network terminate at the MAC Forwarder.

The MSAP interface user **MAY** be the NSI-RFI Forwarding process or the CMTS's host protocol stack.

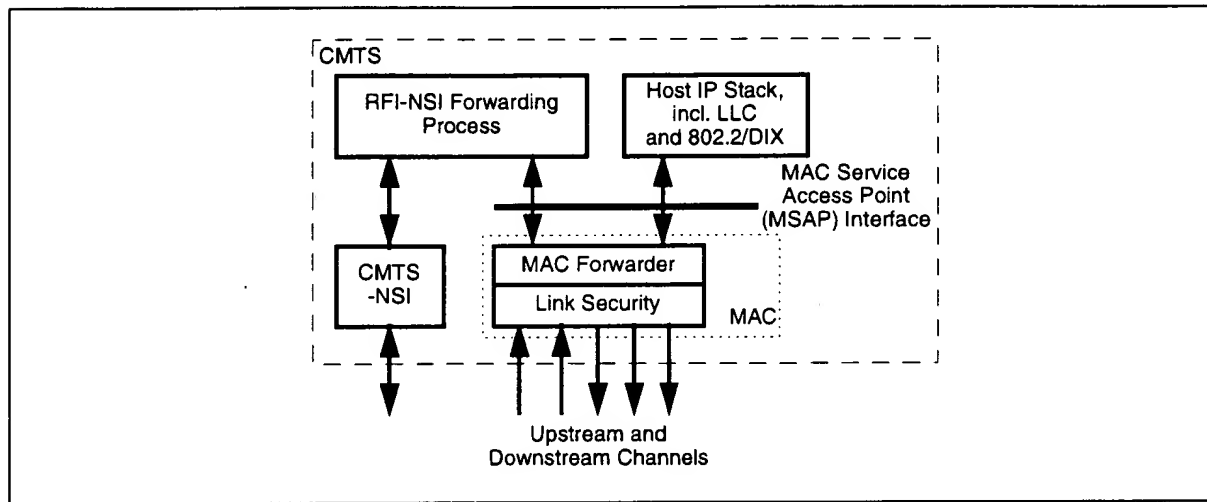


Figure 3-4. MAC Forwarder

Delivery of frames may be based on data-link-layer (bridging) semantics, network-layer (routing) semantics, or some combination. Higher-layer semantics may also be employed (e.g., filters on UDP port numbers). The CMTS **MUST** provide IP connectivity between hosts attached to cable modems, and must do so in a way that meets the expectations of Ethernet-attached customer equipment. For example, the CMTS must either forward ARP packets or it must facilitate a proxy ARP service. The CMTS MAC Forwarder **MAY** provide service for non-IP protocols.

Note that there is no requirement that all upstream and downstream channels be aggregated under one MSAP as shown above. The vendor could just as well choose to implement multiple MSAPs, each with a single upstream and downstream channel.

3.2.1 Rules for Data-Link-Layer Forwarding

If the MAC Forwarder is implemented using only data-link-layer semantics, then the requirements in this section apply.

Delivery of frames is dependent on the Destination Address within the frame. The means of learning the location of each address is vendor-dependent, and MAY include:

- Transparent-bridging-like source-address learning and aging
- Gleaning from MAC Registration Request messages
- Administrative means.

If the destination address of a frame is unicast, and that address is associated with a particular downstream channel, then the frame **MUST** be forwarded to that channel.¹

If the destination address of a frame is unicast, and that address is known to reside on the other (upper) side of the MSAP interface, then the frame **MUST** be delivered to the MSAP interface.

If the destination address is broadcast, multicast², or unknown, the frame **MUST** be delivered to both the MSAP and to all downstream channels. (With the exception of the 3.3.1.1 multicast forwarding rules.)

Delivery rules are similar to those for transparent bridging:

- Frames **MUST NOT** be duplicated.
- Frames that cannot be delivered in a timely fashion **MUST** be discarded.
- The Frame Check Sequence **SHOULD** be preserved rather than regenerated.
- Frames, on a given Service Flow (refer to Section 6.1.2.3), **MUST** be delivered in the order they are received.

3.3 Network Layer

As stated above, the purpose of the data-over-cable system is to transport IP traffic transparently through the system.

The Network Layer protocol is the Internet Protocol (IP) version 4, as defined in [RFC-791], and migrating to IP version 6.

This document imposes no requirements for reassembly of IP packets.

1. Vendors may implement extensions, similar to static addresses in 802.1d/ISO 10038 bridging, that cause such frames to be filtered or handled in some other manner.

2. All multicasts, including 802.1d/ISO 10038 Spanning Tree Bridge BPDU's, **MUST** be forwarded.

3.3.1 Requirements for IGMP Management

3.3.1.1 CMTS Rules

- If link-layer forwarding is used, the CMTS **MUST** forward all Membership Queries on all downstream channels using the appropriate 802.3 multicast group (e.g. 01:00:5E:xx:xx:xx where xx:xx:xx are the low order 23 bits of the multicast address expressed in hex notation. Refer to [IMA].)
- The CMTS **MUST** forward the first copy of Solicited and Unsolicited Membership Reports for any given group received on its upstream RF interface to all of its downstream RF interfaces. However, if membership is managed on a per downstream RF interface basis, Membership Reports and IGMP v2 Leave messages **MAY** be forwarded only on the downstream interface to which the reporting CPE's CM is connected.
- The CMTS **SHOULD** suppress the transmission of additional Membership Reports (for any given group) downstream for at least the Query Response Interval. If the CMTS uses data-link-layer forwarding, it **MUST** also forward the Membership Report out all appropriate Network Side Interfaces.
- The CMTS **SHOULD** suppress the downstream transmission of traffic to any IP multicast group that does not have subscribers on that downstream RF interface (subject to any administrative controls).
- If the CMTS performs network-layer forwarding of multicast packets, it **MUST** implement the router portion of the IGMP protocol [RFC-2236] and **MUST** act as the only IGMP v2 Querier on its downstream RF interfaces.

3.3.1.2 CM Rules

The CM **MUST** support IGMP with the following cable-specific rules. The following requirements apply to conformant CMs:

- The CM **MUST NOT** forward Membership Queries from its CPE interface to its RF interface.
- The CM **MUST NOT** forward Membership Reports or IGMP v2 Leaves received on its RF interface to its CPE interface.
- The CM **MUST NOT** forward multicast traffic from its RF interface to its CPE interface unless a device on its CPE interface is a member of that IP multicast group.
- The CM **MUST** forward multicast traffic from its CPE interface to its RF interface unless administratively (via configuration or other mechanism) prohibited.
- The CM **MUST** forward traffic for the ALL-HOSTS multicast group from its RF interface to its CPE interface unless administratively prohibited. The CPE **MUST** always be considered a member of this group.
- The CM **MUST** forward ALL-HOSTS Group Queries and Group Specific Queries that pass permit filters on its RF interface to its CPE interface or the CM **MUST** implement the Host portion of the IGMP v2 protocol [RFC-2236] on its RF interface for CPEs with active groups and **MUST NOT** act as a Querier on its RF interface. If the CM implements the Host portion of the IGMPv2 protocol, it **MUST** act as an IGMPv2 Querier on its CPE interface. The CM **MUST NOT** require any specific configuration for the associated multicast timer values and **MUST** be capable of adhering to the timers specified in this section. The CM **MAY** provide configuration control that overrides the default values of these timers.
- The CM **MUST** derive the Membership Query Interval by looking at the inter-arrival times of the Membership Query messages. Formally: If $n < 2$, $MQI = 125$ else $MQI = \text{MAX}(125, MQ_n - MQ_{n-1})$, where MQI is the Membership Query Interval in seconds, n is the number of Membership Queries seen, and MQ_n is the epoch time at which the n th Membership Query was seen to the nearest second.
- The Query Response Interval is carried in the Membership Query packet. The Query Response Interval **MUST** be assumed to be 10 seconds if not otherwise set (or set to 0) in the Membership Query packet.

- As a result of receiving a Membership Report on its CPE interface, the CM MUST begin forwarding traffic for the appropriate IP multicast group. The CM MUST stop forwarding multicast traffic from the RF to the CPE side whenever the CM has not received a Membership Report from the CPE side for more than the Membership Interval, which is $(2 * MQI) + QRI$, where MQI is the Membership Query Interval and QRI is the Query Response Interval.
- If the CM has received a Membership Report on its downstream RF interface for groups active on the CM's CPE interface within the Query Response Interval, it MUST suppress transmission on its upstream RF interface of all Membership Reports received on its CPE interface for that group.
- The CM MAY stop forwarding traffic from the RF to the CPE side for a particular multicast group prior to the expiration of the Membership Interval (see above) if it can determine (for example, via an IGMP 'LEAVE' message and the appropriate protocol exchange) that there are no CPE devices subscribed to that particular group.
- The CM MUST treat Unsolicited Membership Reports (IGMP 'JOIN's) from CPE as responses to a Membership Query received on its RF interface. Upon receipt of a JOIN from its CPE interface, the CM MUST start a random timer according to the Host State Diagram, specified in [RFC-2236], and MUST use a Query Response Interval of 10 seconds, as specified above. As specified above, if the CM receives a Membership Report on its RF interface for this group during this random time period, it MUST suppress transmission of this Join on its upstream RF interface. The CM MUST suppress all subsequent Membership Reports for this group until such time as the CM receives a Membership Query (General or Specific to this Group) on its RF interface or a IGMPv2 Leave is received for this group from the CPE interface.

Refer to Appendix L for a state transition diagram example of an approach to these requirements.

Note: Nothing in this section would prohibit the CM from being specifically configured to not forward certain multicast traffic as a matter of network policy.

3.4 Above the Network Layer

The subscribers will be able to use the transparent IP capability as a bearer for higher-layer services. Use of these services will be transparent to the CM.

In addition to the transport of user data, there are several network management and operation capabilities which depend upon the Network Layer. These include:

- SNMP (Simple Network Management Protocol, [RFC-1157]), for network management.
- TFTP (Trivial File Transfer Protocol, [RFC-1350]), a file transfer protocol, for downloading software and configuration information, as modified by TFTP Timeout Interval and Transfer Size Options [RFC-2349].
- DHCP (Dynamic Host Configuration Protocol, [RFC-2131]), a framework for passing configuration information to hosts on a TCP/IP network.
- Time of Day Protocol [RFC-868], to obtain the time of day.

3.5 Data Link Layer

The Data Link Layer is divided into sublayers in accordance with [IEEE802], with the addition of Link-Layer security in accordance with [DOCSIS8]. The sublayers, from the top, are:

- Logical Link Control (LLC) sublayer (Class 1 only)
- Link-Layer Security sublayer
- Media Access Control (MAC) sublayer.

3.5.1 LLC Sublayer

The LLC sublayer **MUST** be provided in accordance with [ISO/IEC10039]. Address resolution **MUST** be used as defined in [RFC-826]. The MAC-to-LLC service definition is specified in [ISO/IEC10039].

3.5.2 Link-Layer Security Sublayer

Link-layer security **MUST** be provided in accordance with [DOCSIS8].

3.5.3 MAC Sublayer

The MAC sublayer defines a single transmitter for each downstream channel - the CMTS. All CMs listen to all frames transmitted on the downstream channel upon which they are registered and accept those where the destinations match the CM itself or CPEs reached via the CMCI port. CMs can communicate with other CMs only through the CMTS.

The upstream channel is characterized by many transmitters (CMs) and one receiver (the CMTS). Time in the upstream channel is slotted, providing for Time Division Multiple Access at regulated time ticks. The CMTS provides the time reference and controls the allowed usage for each interval. Intervals may be granted for transmissions by particular CMs, or for contention by all CMs. CMs may contend to request transmission time. To a limited extent, CMs may also contend to transmit actual data. In both cases, collisions can occur and retries are used.

Section 6 describes the MAC-sublayer messages from the CMTS which direct the behavior of the CMs on the upstream channel, as well as messaging from the CMs to the CMTS.

3.5.3.1 MAC Service Definition

The MAC sublayer service definition is in Appendix E.

3.6 Physical Layer

The Physical (PHY) layer is comprised of two sublayers:

- Transmission Convergence sublayer (present in the downstream direction only)
- Physical Media Dependent (PMD) sublayer.

3.6.1 Downstream Transmission Convergence Sublayer

The Downstream Transmission Convergence sublayer exists in the downstream direction only. It provides an opportunity for additional services over the physical-layer bitstream. These additional services might include, for example, digital video. Definition of any such additional services is beyond the scope of this document.

This sublayer is defined as a continuous series of 188-byte MPEG [ITU-T H.222.0] packets, each consisting of a 4-byte header followed by 184 bytes of payload. The header identifies the payload as belonging to the data-over-cable MAC. Other values of the header may indicate other payloads. The mixture of payloads is arbitrary and controlled by the CMTS.

The Downstream Transmission Convergence sublayer is defined in Section 5 of this document.

3.6.2 PMD Sublayer

The Physical Media Dependent sublayer is defined in Section 4 of this document.

3.6.2.1 Interface Points

Three RF interface points are defined at the PMD sublayer:

- a) Downstream output on the CMTS
- b) Upstream input on the CMTS
- c) Cable in/out at the cable modem.

Separate downstream output and upstream input interfaces on the CMTS are required for compatibility with typical downstream and upstream signal combining and splitting arrangements in headends.

4 Physical Media Dependent Sublayer Specification

4.1 Scope

This specification defines the electrical characteristics and protocol for a cable modem (CM) and cable modem termination system (CMTS). It is the intent of this specification to define an interoperable CM and CMTS such that any implementation of a CM can work with any CMTS. It is not the intent of this specification to imply any specific implementation.

4.2 Upstream

4.2.1 Overview

The upstream Physical Media Dependent (PMD) sublayer uses a FDMA/TDMA burst modulation format, which provides five symbol rates and two modulation formats (QPSK and 16QAM). The modulation format includes pulse shaping for spectral efficiency, is carrier-frequency agile, and has selectable output power level. The PMD sublayer format includes a variable-length modulated burst with precise timing beginning at boundaries spaced at integer multiples of 6.25 μ sec apart (which is 16 symbols at the highest data rate).

Each burst supports a flexible modulation, symbol rate, preamble, randomization of the payload, and programmable FEC encoding.

All of the upstream transmission parameters associated with burst transmission outputs from the CM are configurable by the CMTS via MAC messaging. Many of the parameters are programmable on a burst-by-burst basis.

The PMD sublayer can support a near-continuous mode of transmission, wherein ramp-down of one burst MAY overlap the ramp-up of the following burst, so that the transmitted envelope is never zero. The system timing of the TDMA transmissions from the various CMs MUST provide that the center of the last symbol of one burst and the center of the first symbol of the preamble of an immediately following burst are separated by at least the duration of five symbols. The guard time MUST be greater than or equal to the duration of five symbols plus the maximum timing error. Timing error is contributed by both the CM and CMTS. CM timing performance is specified in Section 4. Maximum timing error and guard time may vary with CMTSs from different vendors.

The upstream modulator is part of the cable modem which interfaces with the cable network. The modulator contains the actual electrical-level modulation function and the digital signal-processing function; the latter provides the FEC, preamble prepend, symbol mapping, and other processing steps. This specification is written with the idea of buffering the bursts in the signal processing portion, and with the signal processing portion (1) accepting the information stream a burst at a time, (2) processing this stream into a complete burst of symbols for the modulator, and (3) feeding the properly-timed burst symbol stream to a memoryless modulator at the exact burst transmit time. The memoryless portion of the modulator only performs pulse shaping and quadrature upconversion.

At the Demodulator, similar to the Modulator, there are two basic functional components: the demodulation function and the signal processing function. Unlike the Modulator, the Demodulator resides in the CMTS and the specification is written with the concept that there will be one demodulation function (not necessarily an actual physical demodulator) for each carrier frequency in use. The demodulation function would receive all bursts on a given frequency.

Note: The unit design approach should be cognizant of the multiple-channel nature of the demodulation and signal processing to be carried out at the headend, and partition/share functionality appropriately to optimally leverage the multi-channel application. A Demodulator design supporting multiple channels in a Demodulator unit may be appropriate.

The demodulation function of the Demodulator accepts a varying-level signal centered around a commanded power level and performs symbol timing and carrier recovery and tracking, burst acquisition, and demodulation. Additionally, the demodulation function provides an estimate of burst timing relative to a reference edge, an estimate of received signal power, an estimate of signal-to-noise ratio, and may engage adaptive equalization to mitigate the effects of a) echoes in the cable plant, b) narrowband ingress and c) group delay. The signal-processing function of the Demodulator performs the inverse processing of the signal-processing function of the Modulator. This includes accepting the demodulated burst data stream and decoding, etc., and possibly multiplexing the data from multiple channels into a single output stream. The signal-processing function also provides the edge-timing reference and gating-enable signal to the demodulators to activate the burst acquisition for each assigned burst slot. The signal-processing function may also provide an indication of successful decoding, decoding error, or fail-to-decode for each codeword and the number of corrected Reed-Solomon symbols in each codeword. For every upstream burst, the CMTS has a prior knowledge of the exact burst length in symbols (see Section 4.2.6, Section 4.2.10.1, and Appendix A.2).

4.2.2 Modulation Formats

The upstream modulator **MUST** provide both QPSK and 16QAM modulation formats.

The upstream demodulator **MUST** support QPSK, 16QAM, or both modulation formats.

4.2.2.1 Modulation Rates

The upstream modulator **MUST** provide QPSK at 160, 320, 640, 1,280, and 2,560 ksym/sec, and 16QAM at 160, 320, 640, 1,280, and 2,560 ksym/sec.

This variety of modulation rates, and flexibility in setting upstream carrier frequencies, permits operators to position carriers in gaps in the pattern of narrowband ingress, as discussed in Appendix G.

The symbol rate for each upstream channel is defined in an Upstream Channel Descriptor (UCD) MAC message. All CM's using that upstream channel **MUST** use the defined symbol rate for upstream transmissions.

4.2.2.2 Symbol Mapping

The modulation mode (QPSK or 16QAM) is programmable. The symbols transmitted in each mode and the mapping of the input bits to the I and Q constellation **MUST** be as defined in Table 4-1. In the table, I_1 is the MSB of the symbol map, Q_1 is the LSB for QPSK, and Q_0 is the LSB for 16QAM. Q_1 and I_0 have intermediate bit positions in 16QAM. The MSB **MUST** be the first bit in the serial data into the symbol mapper.

Table 4-1. I/Q Mapping

QAM Mode	Input bit Definitions
QPSK	$I_1 Q_1$
16QAM	$I_1 Q_1 I_0 Q_0$

The upstream QPSK symbol mapping MUST be as shown in Figure 4-1.




Figure 4-1. QPSK Symbol Mapping

The 16QAM non-inverted (Gray-coded) symbol mapping MUST be as shown in Figure 4-2.




Figure 4-2. 16QAM Gray-Coded Symbol Mapping

The 16QAM differential symbol mapping MUST be as shown in Figure 4-3.




Figure 4-3. 16QAM Differential-Coded Symbol Mapping

If differential quadrant encoding is enabled, then the currently-transmitted symbol quadrant is derived from the previously transmitted symbol quadrant and the current input bits via Table 4-2.

Table 4-2. Derivation of Currently Transmitted Symbol Quadrant

Current Input Bits I(1) Q(1)	Quadrant Phase Change	MSBs of Previously Transmitted Symbol	MSBs for Currently Transmitted Symbol
00	0°	11	11
00	0°	01	01
00	0°	00	00
00	0°	10	10
01	90°	11	01
01	90°	01	00
01	90°	00	10
01	90°	10	11
11	180°	11	00
11	180°	01	10
11	180°	00	11
11	180°	10	01
10	270°	11	10
10	270°	01	11
10	270°	00	01
10	270°	10	00

4.2.2.3 Spectral Shaping

The upstream PMD sublayer MUST support a 25% Nyquist square root raised cosine shaping.

The occupied spectrum MUST NOT exceed the channel widths shown in Table 4-3.

Table 4-3. Maximum Channel Width

Symbol Rate (ksym/sec)	Channel Width (kHz) ^a
160	200
320	400
640	800
1,280	1,600
2,560	3,200

a. Channel width is the -30 dB bandwidth.

4.2.2.4 Upstream Frequency Agility and Range

The upstream PMD sublayer MUST support operation over the frequency range of 5-42 MHz edge to edge.

Offset frequency resolution MUST be supported having a range of ± 32 kHz (increment = 1 Hz; implement within ± 10 Hz).

4.2.2.5 Spectrum Format

The upstream modulator **MUST** provide operation with the format $s(t) = I(t) \cdot \cos(\omega t) - Q(t) \cdot \sin(\omega t)$, where t denotes time and ω denotes angular frequency.

4.2.3 FEC Encode

4.2.3.1 FEC Encode Modes

The upstream modulator **MUST** be able to provide the following selections: Reed-Solomon codes over GF(256) with $T = 1$ to 10 or no FEC coding.

The following Reed-Solomon generator polynomial **MUST** be supported:

$$g(x) = (x + \alpha^0)(x + \alpha^1) \dots (x + \alpha^{2T-1}) \text{ where the primitive element alpha is 0x02 hex}$$

The following Reed-Solomon primitive polynomial **MUST** be supported:

$$p(x) = x^8 + x^4 + x^3 + x^2 + 1$$

The upstream modulator **MUST** provide codewords from a minimum size of 18 bytes (16 information bytes [k] plus two parity bytes for $T = 1$ error correction) to a maximum size of 255 bytes (k -bytes plus parity-bytes). The uncoded word size can have a minimum of one byte.

In Shortened Last Codeword mode, the CM **MUST** provide the last codeword of a burst shortened from the assigned length of k data bytes per codeword as described in Section 4.2.10.1.2 of this document.

The value of T **MUST** be configured in response to the Upstream Channel Descriptor from the CMTS.

4.2.3.2 FEC Bit-to-Symbol Ordering

The input to the Reed-Solomon Encoder is logically a serial bit stream from the MAC layer of the CM, and the first bit of the stream **MUST** be mapped into the MSB of the first Reed-Solomon symbol into the encoder. The MSB of the first symbol out of the encoder **MUST** be mapped into the first bit of the serial bit stream fed to the Scrambler.

[Note that the MAC byte-to-serial upstream convention calls for the byte LSB to be mapped into the first bit of the serial bit stream per Section 6.2.1.3.]

4.2.4 Scrambler (Randomizer)

The upstream modulator **MUST** implement a scrambler (shown in Figure 4-4) where the 15-bit seed value **MUST** be arbitrarily programmable.

At the beginning of each burst, the register is cleared and the seed value is loaded. The seed value **MUST** be used to calculate the scrambler bit which is combined in an XOR with the first bit of data of each burst (which is the MSB of the first symbol following the last symbol of the preamble).

The scrambler seed value **MUST** be configured in response to the Upstream Channel Descriptor from the CMTS.

The polynomial **MUST** be $x^{15} + x^{14} + 1$.

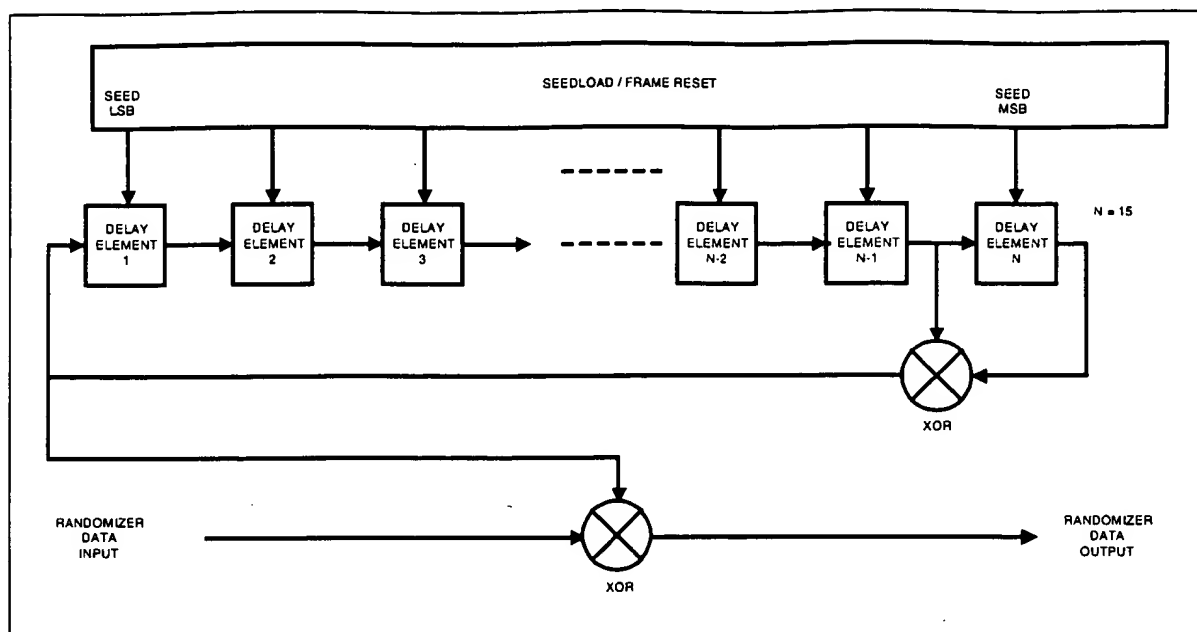


Figure 4-4. Scrambler Structure

4.2.5 Preamble Prepend

The upstream PMD sublayer **MUST** support a variable-length preamble field that is prepended to the data after they have been randomized and Reed-Solomon encoded.

The first bit of the Preamble Pattern is the first bit into the symbol mapper (Figure 4-8), and is I_1 in the first symbol of the burst (see Section 4.2.2.2). The first bit of the Preamble Pattern is designated by the Preamble Value Offset as described in Table 6-19, Section 6.3.3.

The value of the preamble that is prepended **MUST** be programmable and the length **MUST** be 0, 2, 4, ..., or 1024 bits for QPSK and 0, 4, 8, ..., or 1024 bits for 16QAM. Thus, the maximum length of the preamble is 512 QPSK symbols or 256 QAM symbols.

The preamble length and value **MUST** be configured in response to the Upstream Channel Descriptor message transmitted by the CMTS.

4.2.6 Burst Profiles

The transmission characteristics are separated into three portions: a) Channel Parameters, b) Burst Profile Attributes, and c) User Unique Parameters. The Channel Parameters include i) the symbol rate (five rates from 160 ksym/sec to 2.56 Msym/sec in octave steps), ii) the center frequency (Hz), and iii) the 1024-bit Preamble Superstring. The Channel Parameters are further described in Section 6.3.3 Table 6-18; these characteristics are shared by all users on a given channel. The Burst Profile Attributes are listed in Table 4-4, and are further described in Section 6.3.3 Table 6-19; these parameters are the shared attributes corresponding to a burst type. The User Unique Parameters may vary for each user even when using the same burst type on the same channel as another user (for example, Power Level), and are listed in Table 4-5.

The CM **MUST** generate each burst at the appropriate time as conveyed in the mini-slot grants provided by the CMTS MAPs (Section 6.3.4).

The CM MUST support all burst profiles commanded by the CMTS via the Burst Descriptors in the UCD (Section 6.3.3), and subsequently assigned for transmission in a MAP (Section 6.3.4).

Table 4-4. Burst Profile Attributes

Burst Profile Attributes	Configuration Settings
Modulation	QPSK, 16 QAM
Diff Enc	On/Off
Preamble Length	0-1024 bits (Note Section 4.2.5)
Preamble Value offset	0 to 1022
FEC Error Correction (T)	0 to 10 (0 implies no FEC. The number of codeword parity bytes is 2^T)
FEC Codeword Information Bytes (k)	Fixed: 16 to 253 (assuming FEC on) Shortened: 16 to 253 (assuming FEC on)
Scrambler Seed	15 bits
Maximum Burst Length (minislots) ^a	0 to 255
Guard Time	5 to 255 symbols
Last Codeword Length	Fixed, shortened
Scrambler On/Off	On/Off

- a. A burst length of 0 mini-slots in the Channel Profile means that the burst length is variable on that channel for that burst type. The burst length, while not fixed, is granted explicitly by the CMTS to the CM in the MAP.

Table 4-5. User Unique Burst Parameters

User Unique Parameter	Configuration Settings
Power Level ^a	+8 to +55 dBmV (16QAM) +8 to +58 dBmV (QPSK) 1-dB steps
Offset Frequency ^a	Range = ± 32 kHz; increment = 1 Hz; implement within ± 10 Hz
Ranging Offset	0 to $(2^{16} - 1)$, increments of 6.25 $\mu\text{sec}/64$
Burst Length (mini-slots) if variable on this channel (changes burst-to-burst)	1 to 255 mini-slots
Transmit Equalizer Coefficients ^a (advanced modems only)	Up to 64 coefficients; 4 bytes per coefficient: 2 real and 2 complex

- a. Values in table apply for this given channel and symbol rate.

The CM MUST implement the Offset Frequency to within ± 10 Hz.

Ranging Offset is the delay correction applied by the CM to the CMTS Upstream Frame Time derived at the CM. It is an advancement equal to roughly the round-trip delay of the CM from the CMTS, and is needed to synchronize upstream transmissions in the TDMA scheme. The CMTS MUST provide feedback correction for this offset to the CM, based on reception of one or more successfully received bursts (i.e., satisfactory result from each technique employed: error correction and/or CRC), with accuracy within 1/2 symbol and resolution of 1/64 of the frame tick increment ($6.25 \mu\text{sec}/64 = 0.09765625 \mu\text{sec} = 1/4$ the symbol duration of the highest symbol rate = 10.24 MHz^{-1}). The CMTS sends adjustments to the CM, where a negative value implies the Ranging Offset is to be decreased, resulting in later times of transmission at the CM. The CM MUST implement the correction with resolution of at most 1 symbol duration (of the symbol rate in use for a given burst), and (other than a fixed bias) with accuracy within $\pm 0.25 \mu\text{sec}$ plus $\pm 1/2$ symbol owing to resolution. The accuracy of CM burst timing of

$\pm 0.25 \mu\text{sec}$ plus $\pm 1/2$ symbol is relative to the mini-slot boundaries derivable at the CM based on an ideal processing of the timestamp signals received from the CMTS.

The CM **MUST** be capable of switching burst profiles with no reconfiguration time required between bursts except for changes in the following parameters: 1) Output Power, 2) Modulation, 3) Symbol Rate, 4) Offset frequency, 5) Channel Frequency, and 6) Ranging Offset.

For Symbol Rate, Offset frequency and Ranging Offset, the CM **MUST** be able to transmit consecutive bursts as long as the CMTS allocates at least 96 symbols in between the last symbol center of one burst and the first symbol center of the following burst. The maximum reconfiguration time of 96 symbols should compensate for the ramp down time of one burst and the ramp up time of the next burst as well as the overall transmitter delay time including the pipeline delay and optional pre-equalizer delay. For modulation type changes, the CM **MUST** be able to transmit consecutive bursts as long as the CMTS allocates at least 96 symbols in between the last symbol center of one burst and the first symbol center of the following burst. Output Power, Symbol Rate, Offset frequency, Channel Frequency and Ranging Offset **MUST NOT** be changed until the CM is provided sufficient time between bursts by the CMTS. Transmitted Output Power, Symbol Rate, Offset frequency, Channel Frequency and Ranging Offset **MUST NOT** change while more than -30 dB of any symbol's energy of the previous burst remains to be transmitted, or more than -30 dB of any symbol's energy of the next burst has been transmitted. The modulation **MUST NOT** change while more than -30 dB of any symbol's energy of the previous burst remains to be transmitted, or more than -30 dB of any symbol's energy of the next burst has been transmitted, **EXCLUDING** the effect of the transmit equalizer (if present in the CM). [This is to be verified with the transmit equalizer providing no filtering; delay only, if that. Note that if the CMTS has decision feedback in its equalizer, it may need to provide more than the 96 symbol gap between bursts of different modulation type which the same CM may use; this is a CMTS decision.] Negative ranging offset adjustments will cause the 96 symbol guard to be violated. The CMTS must assure that this does not happen by allowing extra guard time between bursts that is at least equal to the amount of negative ranging offset.

If Channel Frequency is to be changed, then the CM **MUST** be able to implement the change between bursts as long as the CMTS allocates at least 96 symbols plus 100 msec between the last symbol center of one burst and the first symbol of the following burst.

The Channel Frequency of the CM **MUST** be settled within the phase noise and accuracy requirements of Section 4.2.9.5 and Section 4.2.9.6 within 100 msec from the beginning of the change.

If Output Power is to be changed by 1 dB or less, the CM **MUST** be able to implement the change between bursts as long as the CMTS allocates at least 96 symbols plus 5 usec between the last symbol center of one burst and the first symbol center of the following burst.

If Output Power is to be changed by more than 1 dB, the CM **MUST** be able to implement the change between bursts as long as the CMTS allocates at least 96 symbols plus 10 usec between the last symbol center of one burst and the first symbol center of the following burst.

The Output Power of the CM **MUST** be settled to within ± 0.1 dB of its final output power level a) within 5 μsec from the beginning of a change of 1 dB or less, and b) within 10 μsec from the beginning of a change of greater than 1 dB.

The output transmit power **MUST** be maintained constant within a TDMA burst to within less than 0.1 dB (excluding the amount theoretically present due to pulse shaping, and amplitude modulation in the case of 16 QAM).

4.2.7 Burst Timing Convention

Figure 4-5 illustrates the nominal burst timing.

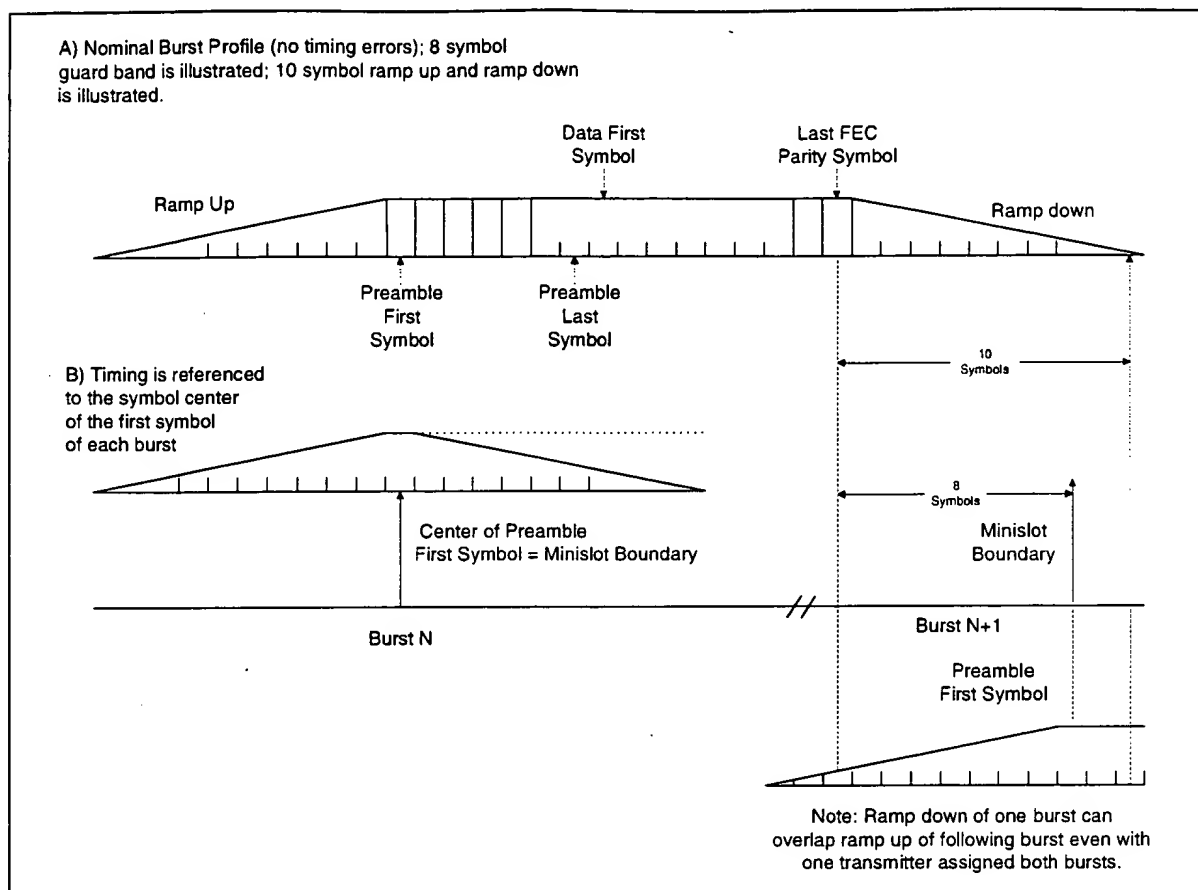


Figure 4-5. Nominal Burst Timing

Figure 4-6 indicates worst-case burst timing. In this example, burst N arrives 1.5 symbols late, and burst N+1 arrives 1.5 symbols early, but separation of 5 symbols is maintained; 8-symbol guardband shown.




Figure 4-6. Worst-Case Burst Timing

At a symbol rate of R_s , symbols occur at a rate of one each $T_s = 1/R_s$ seconds. Ramp Up and Ramp Down are the spread of a symbol in the time domain beyond T_s duration owing to the symbol-shaping filter. If only one symbol were transmitted, its duration would be longer than T_s due to the shaping filter impulse response being longer than T_s . The spread of the first and last symbols of a burst transmission effectively extends the duration of the burst to longer than $N * T_s$, where N is the number of symbols in the burst.

4.2.8 Transmit Power Requirements

The upstream PMD sublayer **MUST** support varying the amount of transmit power. Requirements are presented for 1) the range of commanded transmit power, 2) the step size of the power commands, and 3) the accuracy (actual output power compared to the commanded amount) of the response to the command.

The mechanism by which power adjustments are performed is defined in Section 9.2.4 of this document. Such adjustments **MUST** be within the ranges of tolerances described below.

4.2.8.1 Output Power Agility and Range

The output transmit power in the design bandwidth **MUST** be variable over the range of +8 dBmV to 55 dBmV (16 QAM) or 58 dBmV (QPSK), in 1-dB steps.

The absolute accuracy of the transmitted power **MUST** be ± 2 dB, and the step size accuracy ± 0.4 dB, with an allowance for hysteresis while switching in/out a step attenuator (e.g. 20 dB) in which case the accuracy requirement is relaxed to ± 1.4 dB. For example, the actual power increase resulting from a command to increase the power level by 1 dB in a CM's next transmitted burst **MUST** be between 0.6 and 1.4 dB.

The step resolution **MUST** be 1 dB or less. When a CM is commanded with finer resolution than it can implement, it **MUST** round to the nearest supported step size. If the commanded step is half way between two supported step sizes, the CM **MUST** choose the smaller step. For example, with a supported step resolution of 1 dB, a command to step ± 0.5 dB would result in no step, while a command to step ± 0.75 dB would result in a ± 1 dB step.

4.2.9 Fidelity Requirements

4.2.9.1 Spurious Emissions

The noise and spurious power **MUST NOT** exceed the levels given in Table 4-6, Table 4-7, and Table 4-8.

In Table 4-6, Inband spurious includes noise, carrier leakage, clock lines, synthesizer spurious products, and other undesired transmitter products. It does not include ISI. The measurement bandwidth for Inband spurious is equal to the symbol rate (e.g., 160 kHz for 160 ksym/sec).

The measurement bandwidth for the 3 (or fewer) Carrier-Related Frequency Bands (below 42 MHz) is 160 kHz, with up to three 160 kHz bands, each with no more than -47 dBc, allowed to be excluded from the "Bands within 5 to 42 MHz Transmitting Burst" specs of Table 4-8.

The measurement bandwidth is also 160 kHz for the Between Bursts specs of Table 4-6 below 42 MHz; the Transmitting Burst specs apply during the mini-slots granted to the CM (when the CM uses all or a portion of the grant), and for a minislot before and after the granted mini-slots. [Note that a minislot may be as short as 32 symbols, or 12.5 microseconds at the 2.56 Msym/sec rate, or as short as 200 microseconds at the 160 ksym/sec rate.] The Between Bursts specs apply except during a used grant of mini-slots, and the minislot before and after the used grant.

Table 4-6. Spurious Emissions

Parameter	Transmitting Burst	Between Bursts
Inband (Inband spurious includes noise, carrier leakage, clock lines, synthesizer spurious products, and other undesired transmitter products. It does not include Inter Symbol Interference (ISI)).	-40 dBc	The greater of -72 dBc or -59 dBmV
Adjacent Band	See Table 4-7	The greater of -72 dBc or -59 dBmV
3 or Fewer Carrier-Related Frequency Bands(such as second harmonic, if < 42 MHz)	-47 dBc	The greater of -72 dBc or -59 dBmV
Bands within 5 to 42 MHz(excluding assigned channel, adjacent channels, and carrier-related channels)	See Table 4-8	The greater of -72 dBc or -59 dBmV
CM Integrated Spurious Emissions Limits (all in 4 MHz, includes discretes) ^a		
42 to 54 MHz	max(-40 dBc, -26 dBmV)	-26 dBmV
54 to 60 MHz	-35 dBmV	-40 dBmV
60 to 88 MHz	-40 dBmV	-40 dBmV
88-860 MHz	-45 dBmV	max(-45 dBmV, -40 dBc ^b)
CM Discrete Spurious Emissions Limits ^a		
42 to 54 MHz	-max(-50 dBc, -36 dBmV)	-36 dBmV
54 to 88 MHz	-50 dBmV	-50 dBmV
88 to 860 MHz	-50 dBmV	-50 dBmV

- These spec limits exclude a single discrete spur related to the tuned received channel; this single discrete spur MUST be no greater than -40 dBmV.
- "dBc" is relative to the received downstream signal level. Some spurious outputs are proportional to the receive signal level.

4.2.9.1.1 Adjacent Channel Spurious Emissions

Spurious emissions from a transmitted carrier may occur in an adjacent channel which could be occupied by a carrier of the same or different symbol rates. The following table lists the required adjacent channel spurious emission levels for all combinations of transmitted carrier symbol rates and adjacent channel symbol rates. The measurement is performed in an adjacent channel interval that is of appropriate bandwidth and distance from the transmitted carrier based on the symbol rates of the transmitted carrier and the carrier in the adjacent channel.

Table 4-7. Adjacent Channel Spurious Emissions Relative to the Transmitted Burst Power Level

Transmitted carrier symbol rate	Specification in the interval	Measurement interval and distance from carrier edge	Adjacent channel carrier symbol rate
160 Ksym/sec	-45 dBc	20 KHz to 180 KHz	160 Ksym/sec
	-45 dBc	40 KHz to 360 KHz	320 Ksym/sec
	-45 dBc	80 KHz to 720 KHz	640 Ksym/sec
	-42 dBc	160 KHz to 1440 KHz	1280 Ksym/sec
	-39 dBc	320 KHz to 2880 KHz	2560 Ksym/sec
All other symbol rates	-45 dBc	20 KHz to 180 KHz	160 Ksym/sec
	-45 dBc	40 KHz to 360 KHz	320 Ksym/sec
	-45 dBc	80 KHz to 720 KHz	640 Ksym/sec
	-44 dBc	160 KHz to 1440 KHz	1280 Ksym/sec
	-41 dBc	320 KHz to 2880 KHz	2560 Ksym/sec

4.2.9.1.2 Spurious Emissions in 5 to 42 MHz

Spurious emissions, other than those in an adjacent channel or carrier related emissions listed above, may occur in intervals that could be occupied by other carriers of the same or different symbol rates. To accommodate these different symbol rates and associated bandwidths, the spurious emissions are measured in an interval equal to the bandwidth corresponding to the symbol rate of the carrier that could be transmitted in that interval. This interval is independent of the current transmitted symbol rate.

The following table lists the possible symbol rates that could be transmitted in an interval, the required spurious level in that interval, and the initial measurement interval at which to start measuring the spurious emissions. Measurements should start at the initial distance and be repeated at increasing distance from the carrier until the upstream band edge, 5 MHz or 42 MHz, is reached. Measurement intervals should not include carrier related emissions.

Table 4-8. Spurious Emissions in 5 to 42 MHz Relative to the Transmitted Burst Power Level

Possible symbol rate in this interval	Specification in the interval	Initial measurement interval and distance from carrier edge
160 Ksym/sec	-53 dBc	220 KHz to 380 KHz
320 Ksym/sec	-50 dBc	240 KHz to 560 KHz
640 Ksym/sec	-47 dBc	280 KHz to 920 KHz
1280 Ksym/sec	-44 dBc	360 KHz to 1640 KHz
2560 Ksym/sec	-41 dBc	520 KHz to 3080 KHz

4.2.9.2 Spurious Emissions During Burst On/Off Transients

Each transmitter **MUST** control spurious emissions, prior to and during ramp up and during and following ramp down, before and after a burst in the TDMA scheme.

On/off spurious emissions, such as the change in voltage at the upstream transmitter output due to enabling or disabling transmission, **MUST** be no more than 100 mV, and such a step **MUST** be dissipated no faster than 2 μ s of constant slewing. This requirement applies when the CM is transmitting at +55 dBmV or more; at backed-off transmit levels, the maximum change in voltage **MUST** decrease by a factor of 2 for each 6-dB decrease of power level from +55 dBmV, down to a maximum change of 7 mV at 31 dBmV and below. This requirement does not apply to CM power-on and power-off transients.

4.2.9.3 Symbol Error Rate (SER)

Modulator performance **MUST** be within 0.5 dB of theoretical SER vs C/N (i.e., E_s/N_0), for SER as low as 10^{-6} uncoded, for QPSK and 16 QAM.

The SER degradation is determined by the cluster variance caused by the transmit waveform at the output of an ideal square-root raised-cosine receive filter. It includes the effects of ISI, spurious, phase noise, and all other transmitter degradations.

Cluster SNR should be measured on a modulation analyzer using a square-root raised cosine receive filter with $\alpha = 0.25$. The measured SNR **MUST** be better than 30 dB.

4.2.9.4 Filter Distortion

The following requirements assume that any pre-equalization is disabled.

4.2.9.4.1 Amplitude

The spectral mask **MUST** be the ideal square root raised cosine spectrum with $\alpha = 0.25$, within the ranges given below:

$$f_c - R_s/4 \text{ Hz to } f_c + R_s/4 \text{ Hz: } -0.3 \text{ dB to } +0.3 \text{ dB}$$

$$f_c - 3R_s/8 \text{ Hz to } f_c - R_s/4 \text{ Hz, and } f_c + R_s/4 \text{ Hz to } f_c + 3R_s/8 \text{ Hz: } -0.5 \text{ dB to } 0.3 \text{ dB}$$

$$f_c - R_s/2 \text{ Hz and } f_c + R_s/2 \text{ Hz: } -3.5 \text{ dB to } -2.5 \text{ dB}$$

$$f_c - 5R_s/8 \text{ Hz and } f_c + 5R_s/8 \text{ Hz: no greater than } -30 \text{ dB}$$

where f_c is the center frequency, R_s is the symbol rate, and the spectral density is measured with a resolution bandwidth of 10 KHz or less.

4.2.9.4.2 Phase

$$f_c - 5R_s/8 \text{ Hz to } f_c + 5R_s/8 \text{ Hz: Group Delay Variation MUST NOT be greater than } 100 \text{ nsec.}$$

4.2.9.5 Carrier Phase Noise

The upstream transmitter total integrated phase noise (including discrete spurious noise) **MUST** be less than or equal to -43 dBc summed over the spectral regions spanning 1 kHz to 1.6 MHz above and below the carrier.

4.2.9.6 Channel Frequency Accuracy

The CM **MUST** implement the assigned channel frequency within ± 50 parts per million over a temperature range of 0 to 40 degrees C up to five years from date of manufacture.

4.2.9.7 Symbol Rate Accuracy

The upstream modulator **MUST** provide an absolute accuracy of symbol rates ± 50 parts per million over a temperature range of 0 to 40 degrees C up to five years from date of manufacture.

4.2.9.8 Symbol Timing Jitter

Peak-to-peak symbol jitter, referenced to the previous symbol zero-crossing, of the transmitted waveform, **MUST** be less than 0.02 of the nominal symbol duration over a 2-sec period. In other words, the difference between the maximum and the minimum symbol duration during the 2-sec period shall be less than 0.02 of the nominal symbol duration for each of the five upstream symbol rates.

The peak-to-peak cumulative phase error, referenced to the first symbol time and with any fixed symbol frequency offset factored out, **MUST** be less than 0.04 of the nominal symbol duration over a 0.1-sec period. In other words, the difference between the maximum and the minimum cumulative phase error during the 0.1-sec period shall be less than 0.04 of the nominal symbol duration for each of the five upstream symbol rates. Factoring out a fixed symbol frequency offset is to be done by using the computed mean symbol duration during the 0.1 sec.

4.2.10 Frame Structure

Figure 4-7 shows two examples of the frame structure: one where the packet length equals the number of information bytes in a codeword, and another where the packet length is longer than the number of information bytes in one codeword, but less than in two codewords. Example 1 illustrates the fixed codeword-length mode, and Example 2 illustrates the shortened last codeword mode. These modes are defined in Section 4.2.10.1.

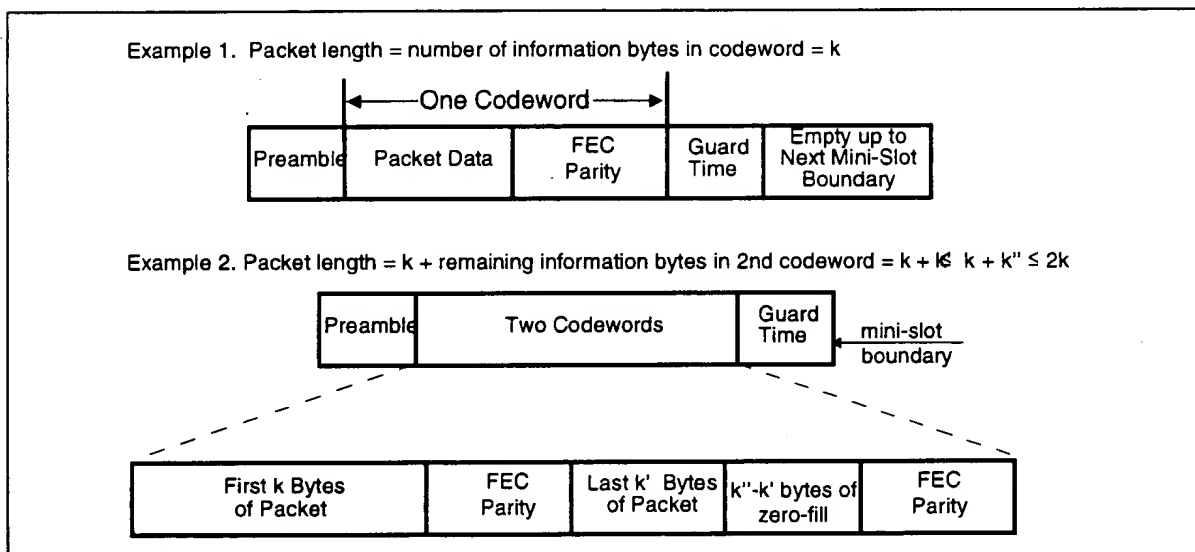


Figure 4-7. Example Frame Structures with Flexible Burst Length Mode

4.2.10.1 Codeword Length

When FEC is enabled, the CM operates in either fix-length codeword mode or in shortened-last codeword mode. The minimum number of information bytes in a codeword in either mode is 16. Shortened-last codeword mode only provides a benefit when the number of bytes in a codeword is greater than the minimum of 16 bytes.

The following descriptions apply to an allocated grant of mini-slots in both contention and non-contention regions. (Allocation of mini-slots is discussed in Section 6 of this document.) The intent of the description is to define rules and conventions such that CMs request the proper number of mini-slots and the CMTS PHY knows what to expect regarding the FEC framing in both fixed codeword length and shortened last codeword modes.

4.2.10.1.1 Fixed Codeword Length

With the fixed-length codewords, after all the data are encoded, zero-fill will occur in this codeword if necessary to reach the assigned k data bytes per codeword, and zero-fill **MUST** continue up to the point when no additional fixed-length codewords can be inserted before the end of the last allocated mini-slot in the grant, accounting for FEC parity and guard-time symbols.

4.2.10.1.2 Shortened Last Codeword

As shown in Figure 4-7, let k' = the number of information bytes that remain after partitioning the information bytes of the burst into full-length (k burst data bytes) codewords. The value of k' is less than k . Given operation in a shortened last codeword mode, let k'' = the number of burst data bytes plus zero-fill bytes in the shortened last codeword. In shortened codeword mode, the CM will encode the data bytes of the burst (including MAC Header) using the assigned codeword size (k information bytes per codeword) until 1) all the data are encoded, or 2) a remainder of data bytes is left over which is less than k . Shortened last codewords shall not have less than 16 information bytes, and this is to be considered when CMs make requests of mini-slots. In shortened last codeword mode, the CM will zero-fill data if necessary until the end of the mini-slot allocation, which in most cases will be the next mini-slot boundary, accounting for FEC parity and guard-time symbols. In many cases, only $k'' - k'$ zero-fill bytes are necessary to fill out a mini-slot allocation with $16 \leq k'' \leq k$ and $k' \leq k''$. However, note the following.

More generally, the CM is required to zero-fill data until the point when no additional fixed-length codewords can be inserted before the end of the last allocated mini-slot in the grant (accounting for FEC parity and guard-time symbols), and then, if possible, a shortened last codeword of zero-fill shall be inserted to fit into the mini-slot allocation.

If, after zero-fill of additional codewords with k information bytes, there are less than 16 bytes remaining in the allocated grant of mini-slots, accounting for parity and guard-time symbols, then the CM shall not create this last shortened codeword.

4.2.11 Signal Processing Requirements

The signal processing order for each burst packet type **MUST** be compatible with the sequence shown in Figure 4-8 and **MUST** follow the order of steps in Figure 4-9.

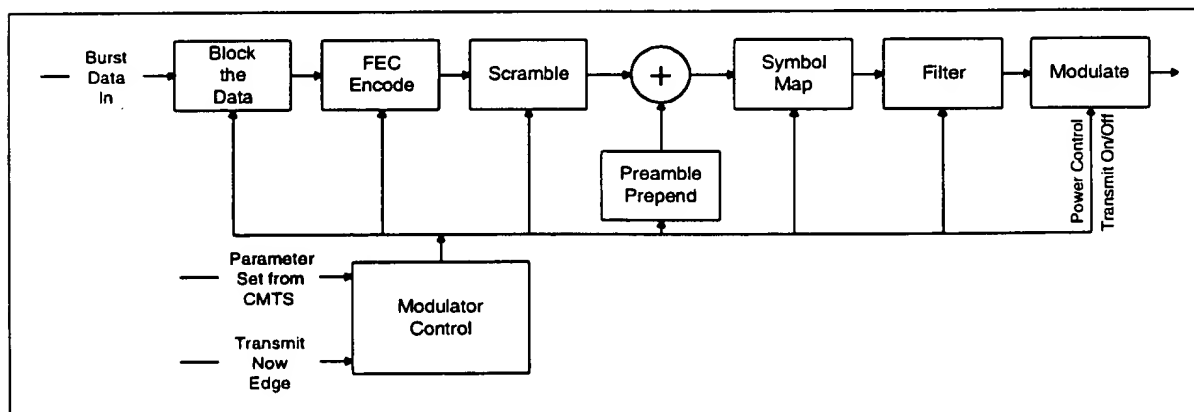


Figure 4-8. Signal-Processing Sequence

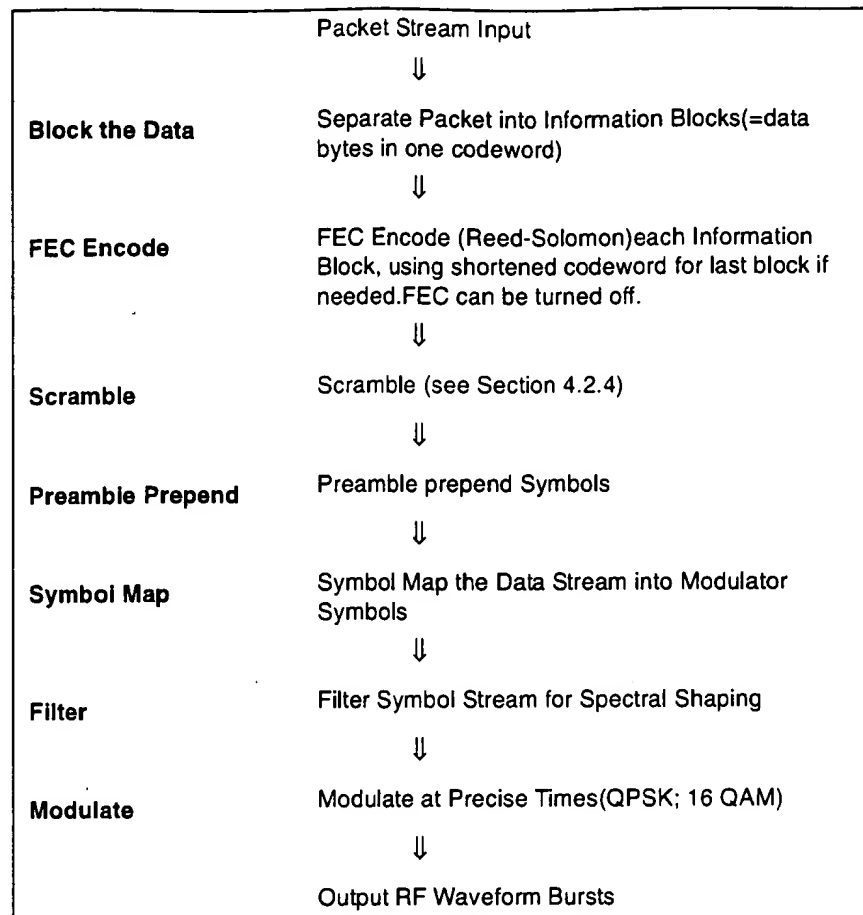


Figure 4-9. TDMA Upstream Transmission Processing

4.2.12 Upstream Demodulator Input Power Characteristics

The maximum total input power to the upstream demodulator MUST NOT exceed 35 dBmV in the 5-42 MHz frequency range of operation.

The intended received power in each carrier MUST be within the values shown in Table 4-9.

Table 4-9. Maximum Range of Commanded Nominal Receive Power in Each Carrier

Symbol Rate (ksym/sec)	Maximum Range (dBmV)
160	-16 to +14
320	-13 to +17
640	-10 to +20
1,280	-7 to +23
2,560	-4 to +26

The demodulator MUST operate within its defined performance specifications with received bursts within ± 6 dB of the nominal commanded received power.

4.2.13 Upstream Electrical Output from the CM

The CM MUST output an RF modulated signal with the characteristics delineated in Table 4-10.

Table 4-10. Electrical Output from CM

Parameter	Value
Frequency	5 to 42 MHz edge to edge
Level range (one channel)	+8 to +55 dBmV (16QAM) +8 to +58 dBmV (QPSK)
Modulation Type	QPSK and 16QAM
Symbol Rate (nominal)	160, 320, 640, 1,280 and 2,560 ksym/sec
Bandwidth	200, 400, 800, 1,600 and 3,200 kHz
Output impedance	75 ohms
Output Return Loss	> 6 dB (5-42 MHz)
Connector	F connector per [IPS-SP-406] (common with the input)

4.3 Downstream

4.3.1 Downstream Protocol

The downstream PMD sublayer MUST conform to ITU-T Recommendations J.83, Annex B for Low-Delay Video Applications [ITU J.83-B], with the exceptions called out in Section 4.3.2.

4.3.2 Scalable Interleaving to Support Low Latency

The downstream PMD sublayer MUST support a variable-depth interleaver with the characteristics defined in Table 4-11. The table contains a subset of interleaver modes found in [ITU J.83-B].

Table 4-11. Interleaver Characteristics

I (Number of Taps)	J (Increment)	Burst Protection 64QAM/256QAM	Latency 64QAM/256QAM
8	16	5.9 μ sec/4.1 μ sec	0.22 msec/0.15 msec
16	8	12 μ sec/8.2 μ sec	0.48 msec/0.33 msec
32	4	24 μ sec/16 μ sec	0.98 msec/0.68 msec
64	2	47 μ sec/33 μ sec	2.0 msec/1.4 msec
128	1	95 μ sec/66 μ sec	4.0 msec/2.8 msec

The interleaver depth, which is coded in a 4-bit control word contained in the FEC frame synchronization trailer, always reflects the interleaving in the immediately-following frame. In addition, errors are allowed while the interleaver memory is flushed after a change in interleaving is indicated.

Refer to [ITU J.83-B] for the control bit specifications required to specify which interleaving mode is used.

4.3.3 Downstream Frequency Plan

The downstream frequency plan should comply with Harmonic Related Carrier (HRC), Incremental Related Carrier (IRC) or Standard (STD) North American frequency plans per [EIA-S542]. However, operation below a center frequency of 91 MHz is not required.

4.3.4 CMTS Output Electrical

The CMTS MUST output an RF modulated signal with the following characteristics defined in Table 4-12.

Table 4-12. CMTS Output

Parameter	Value
Center Frequency (fc)	91 to 857 MHz ± 30 kHz ^a
Level	Adjustable over the range 50 to 61 dBmV
Modulation Type	64QAM and 256QAM
Symbol Rate (nominal)	
64QAM	5.056941 Msym/sec
256QAM	5.360537 Msym/sec
Nominal Channel Spacing	6 MHz
Frequency response	
64QAM	~18% Square Root Raised Cosine shaping
256QAM	~12% Square Root Raised Cosine shaping
Total Discrete Spurious Inband (fc ± 3 MHz)	< -57dBc
Inband Spurious and Noise (fc ± 3 MHz)	< -48dBc; where channel spurious and noise includes all discrete spurious, noise, carrier leakage, clock lines, synthesizer products, and other undesired transmitter products. Noise within ± 50 kHz of the carrier is excluded.
Adjacent channel (fc ± 3.0 MHz) to (fc ± 3.75 MHz)	< -58 dBc in 750 kHz
Adjacent channel (fc ± 3.75 MHz) to (fc ± 9 MHz)	< -62 dBc, in 5.25 MHz, excluding up to 3 spurs, each of which must be < -60 dBc when measured in a 10 kHz band
Next adjacent channel (fc ± 9 MHz) to (fc ± 15 MHz)	Less than the greater of -65 dBc or -12dBmV in 6MHz, excluding up to three discrete spurs. The total power in the spurs must be < -60dBc when each is measured with 10 kHz bandwidth.
Other channels (47 MHz to 1,000 MHz)	< -12dBmV in each 6 MHz channel, excluding up to three discrete spurs. The total power in the spurs must be < -60dBc when each is measured with 10kHz bandwidth.
Phase Noise	1 kHz - 10 kHz: -33dBc double sided noise power 10 kHz - 50 kHz: -51dBc double sided noise power 50 kHz - 3 MHz: -51dBc double sided noise power
Output Impedance	75 ohms
Output Return Loss	> 14 dB within an output channel up to 750 MHz; > 13 dB in an output channel above 750 MHz
Connector	F connector per [IPS-SP-406]

- a. ± 30 kHz includes an allowance of 25 kHz for the largest FCC frequency offset normally built into upconverters.

4.3.5 Downstream Electrical Input to CM

The CM MUST accept an RF modulated signal with the following characteristics (Table 4-13.)

Table 4-13. Electrical Input to CM

Parameter	Value
Center Frequency	91 to 857 MHz \pm 30 kHz
Level Range (one channel)	-15 dBmV to +15 dBmV
Modulation Type	64QAM and 256QAM
Symbol Rate (nominal)	5.056941 Msym/sec (64QAM) and 5.360537 Msym/sec (256QAM)
Bandwidth	6 MHz (18% Square Root Raised Cosine shaping for 64QAM and 12% Square Root Raised Cosine shaping for 256QAM)
Total Input Power (40-900 MHz)	<30 dBmV
Input (load) Impedance	75 ohms
Input Return Loss	> 6 dB (88-860 MHz)
Connector	F connector per [IPS-SP-406] (common with the output)

4.3.6 CM BER Performance

The bit-error-rate performance of a CM MUST be as described in this section. The requirements apply to the I = 128, J = 1 mode of interleaving.

4.3.6.1 64QAM

4.3.6.1.1 64QAM CM BER Performance

Implementation loss of the CM MUST be such that the CM achieves a post-FEC BER less than or equal to 10^{-8} when operating at a carrier to noise ratio (E_s/N_0) of 23.5 dB or greater.

4.3.6.1.2 64QAM Image Rejection Performance

Performance as described in Section 4.3.6.1.1 MUST be met with analog or digital signal at +10 dBc in any portion of the RF band other than the adjacent channels.

4.3.6.1.3 64QAM Adjacent Channel Performance

Performance as described in Section 4.3.6.1.1 MUST be met with a digital signal at 0 dBc in the adjacent channels.

Performance as described in Section 4.3.6.1.1 MUST be met with an analog signal at +10 dBc in the adjacent channels.

Performance as described in Section 4.3.6.1.1, with an additional 0.2-dB allowance, MUST be met with a digital signal at +10 dBc in the adjacent channels.

4.3.6.2 256QAM

4.3.6.2.1 256QAM CM BER Performance

Implementation loss of the CM MUST be such that the CM achieves a post-FEC BER less than or equal to 10^{-8} when operating at a carrier to noise ratio (E_s/N_o) as shown below.

Input Receive Signal Level	E_s/N_o
-6 dBmV to +15dBmV	30dB or greater
Less than -6dBmV down to -15dBmV	33dB or greater

4.3.6.2.2 256QAM Image Rejection Performance

Performance as described in Section 4.3.6.2.1 MUST be met with an analog or a digital signal at +10 dBc in any portion of the RF band other than the adjacent channels.

4.3.6.2.3 256QAM Adjacent Channel Performance

Performance as described in Section 4.3.6.2.1 MUST be met with an analog or a digital signal at 0 dBc in the adjacent channels.

Performance as described in Section 4.3.6.2.1, with an additional 0.5-dB allowance, MUST be met with an analog signal at +10 dBc in the adjacent channels.

Performance as described in Section 4.3.6.2.1, with an additional 1.0-dB allowance, MUST be met with a digital signal at +10 dBc in the adjacent channels.

4.3.7 CMTS Timestamp Jitter

The CMTS timestamp jitter must be less than 500 ns peak-to-peak at the output of the Downstream Transmission Convergence Sublayer. This jitter is relative to an ideal Downstream Transmission Convergence Sublayer that transfers the MPEG packet data to the Downstream Physical Media Dependent Sublayer with a perfectly continuous and smooth clock at the MPEG packet data rate. Downstream Physical Media Dependent Sublayer processing MUST NOT be considered in timestamp generation and transfer to the Downstream Physical Media Dependent Sublayer.

Thus, any two timestamps N_1 and N_2 ($N_2 > N_1$) which were transferred to the Downstream Physical Media Dependent Sublayer at times T_1 and T_2 respectively must satisfy the following relationship:

$$|(N_2 - N_1)/10240000 - (T_2 - T_1)| < 500 \text{ nsec}$$

The jitter includes inaccuracy in timestamp value and the jitter in all clocks. The 500ns allocated for jitter at the Downstream Transmission Convergence Sublayer output must be reduced by any jitter that is introduced by the Downstream Physical Media Dependent Sublayer.

The CM is expected to meet the burst timing accuracy requirements in Section 4.2.6 when the time stamps contain this worst-case jitter.

Note: Jitter is the error (i.e., measured) relative to the CMTS Master Clock. (The CMTS Master Clock is the 10.24 MHz clock used for generating the timestamps.)

The CMTS 10.24 MHz Master Clock MUST have frequency accuracy of $\leq \pm 5$ ppm, drift rate $\leq 10^{-8}$ per second, and edge jitter of ≤ 10 nsec peak-to-peak (± 5 nsec) over a temperature range of 0 to 40degrees C up to ten years from date of manufacture.¹ [The drift rate and jitter requirements on the CMTS Master Clock implies that the duration of two adjacent segments of 10,240,000 cycles will be within 30 nsec, due to 10 nsec jitter on each segments' duration, and 10 nsec due to frequency drift. Durations of other counter lengths also may be deduced: adjacent 1,024,000 segments, ≤ 21 nsec; 1,024,000 length segments separated by one 10,240,000 cycle segment, ≤ 30 nsec; adjacent 102,400,000 segments, ≤ 120 nsec. The CMTS Master Clock MUST meet such test limits in 99% or more measurements.]

-
1. This specification MAY also be met by synchronizing the CMTS Master Clock oscillator to an external frequency reference source. If this approach is used, the internal CMTS Master Clock MUST have frequency accuracy of ± 20 ppm over a temperature range of 0 to 40 degrees C up to 10 years from date of manufacture when no frequency reference source is connected. The drift rate and edge jitter MUST be as specified above.

5 Downstream Transmission Convergence Sublayer

5.1 Introduction

In order to improve demodulation robustness, facilitate common receiving hardware for both video and data, and provide an opportunity for the possible future multiplexing of video and data over the PMD sublayer bitstream defined in Section 4, a sublayer is interposed between the downstream PMD sublayer and the Data-Over-Cable MAC sublayer.

The downstream bitstream is defined as a continuous series of 188-byte MPEG [ITU-T H.222.0] packets. These packets consist of a 4-byte header followed by 184 bytes of payload. The header identifies the payload as belonging to the Data-Over-Cable MAC. Other values of the header may indicate other payloads. The mixture of MAC payloads and those of other services is optional and is controlled by the CMTS.

Figure 5-1 illustrates the interleaving of Data-Over-Cable (DOC) MAC bytes with other digital information (digital video in the example shown).

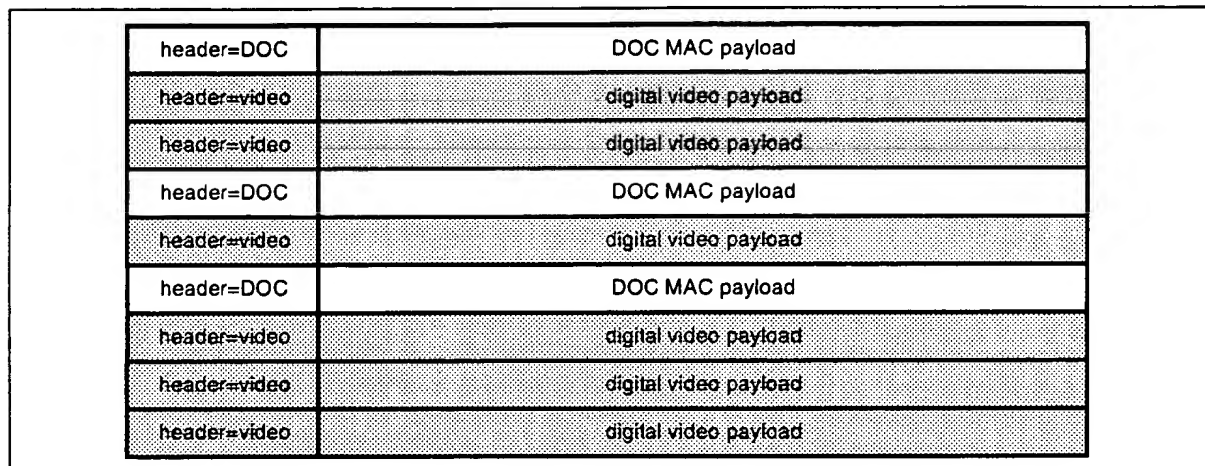


Figure 5-1. Example of Interleaving MPEG Packets in Downstream

5.2 MPEG Packet Format

The format of an MPEG Packet carrying DOCSIS data is shown in Figure 5-2. The packet consists of a 4-byte MPEG Header, a pointer_field (not present in all packets) and the DOCSIS Payload.




Figure 5-2. Format of an MPEG Packet

5.3 MPEG Header for DOCSIS Data-Over-Cable

The format of the MPEG Transport Stream header is defined in Section 2.4 of [ITU-T H.222.0]. The particular field values that distinguish Data-Over-Cable MAC streams are defined in Table 5-1. Field names are from the ITU specification.

The MPEG Header consists of 4 bytes that begin the 188-byte MPEG Packet. The format of the header for use on an DOCSIS Data-Over-Cable PID is restricted to that shown in Table 5-1. The header format conforms to the MPEG standard, but its use is restricted in this specification to NOT ALLOW inclusion of an adaptation_field in the MPEG packets.

Table 5-1. MPEG Header Format for DOCSIS Data-Over-Cable Packets

Field	Length (bits)	Description
sync_byte	8	0x47; MPEG Packet Sync byte
transport_error_indicator	1	Indicates an error has occurred in the reception of the packet. This bit is reset to zero by the sender, and set to one whenever an error occurs in transmission of the packet
payload_unit_start_indicator	1	A value of one indicates the presence of a pointer_field as the first byte of the payload (fifth byte of the packet)
transport_priority	1	Reserved; set to zero
PID (see Note)	13	DOCSIS Data-Over-Cable well-known PID (0x1FFE)
transport_scrambling_control	2	Reserved, set to '00'
adaptation_field_control	2	'01'; use of the adaptation_field is NOT ALLOWED on the DOCSIS PID
continuity_counter	4	cyclic counter within this PID

5.4 MPEG Payload for DOCSIS Data-Over-Cable

The MPEG payload portion of the MPEG packet will carry the DOCSIS MAC frames. The first byte of the MPEG payload will be a 'pointer_field' if the payload_unit_start_indicator (PUSI) of the MPEG header is set.

stuff_byte

This standard defines a stuff_byte pattern having a value (0xFF) that is used within the DOCSIS payload to fill any gaps between the DOCSIS MAC frames. This value is chosen as an unused value for the first byte of the DOCSIS MAC frame. The 'FC' byte of the MAC Header will be defined to never contain this value. (FC_TYPE = '11' indicates a MAC-specific frame, and FC_PARM = '11111' is not currently used and, according to this specification, is defined as an illegal value for FC_PARM.)

pointer_field

The pointer_field is present as the fifth byte of the MPEG packet (first byte following the MPEG header) whenever the PUSI is set to one in the MPEG header. The interpretation of the pointer_field is as follows:

The pointer_field contains the number of bytes in this packet that immediately follow the pointer_field that the CM decoder must skip past before looking for the beginning of an DOCSIS MAC Frame. A pointer field MUST be present if it is possible to begin a Data-Over-Cable MAC Frame in the packet, and MUST point to either:

1. the beginning of the first MAC frame to start in the packet or
2. to any stuff_byte preceding the MAC frame.

5.5 Interaction with the MAC Sublayer

MAC frames may begin anywhere within an MPEG packet, MAC frames may span MPEG packets, and several MAC frames may exist within an MPEG packet.

The following figures show the format of the MPEG packets that carry DOCSIS MAC frames. In all cases, the PUSI flag indicates the presence of the pointer_field as the first byte of the MPEG payload.

Figure 5-3 shows a MAC frame that is positioned immediately after the pointer_field byte. In this case, pointer_field is zero, and the DOCSIS decoder will begin searching for a valid FC byte at the byte immediately following the pointer_field.




Figure 5-3. Packet Format Where a MAC Frame Immediately Follows the pointer_field

Figure 5-4 shows the more general case where a MAC Frame is preceded by the tail of a previous MAC Frame and a sequence of stuffing bytes. In this case, the pointer_field still identifies the first byte after the tail of Frame #1 (a stuff_byte) as the position where the decoder should begin searching for a legal MAC sublayer FC value. This format allows the multiplexing operation in the CMTS to immediately insert a MAC frame that is available for transmission if that frame arrives after the MPEG header and pointer_field have been transmitted.

In order to facilitate multiplexing of the MPEG packet stream carrying DOCSIS data with other MPEG-encoded data, the CMTS SHOULD NOT transmit MPEG packets with the DOCSIS PID which contain only stuff_bytes in the payload area. MPEG null packets SHOULD be transmitted instead. Note that there are timing relationships implicit in the DOCSIS MAC sublayer which must also be preserved by any MPEG multiplexing operation.




Figure 5-4. Packet Format with MAC Frame Preceded by Stuffing Bytes

Figure 5-5 shows that multiple MAC frames may be contained within the MPEG packet. The MAC frames may be concatenated one after the other or be separated by an optional sequence of stuffing bytes.




Figure 5-5. Packet Format Showing Multiple MAC Frames in a Single Packet

Figure 5-6 shows the case where a MAC frame spans multiple MPEG packets. In this case, the pointer_field of the succeeding frame points to the byte following the last byte of the tail of the first frame.

MPEG Header (PUSI = 1)	pointer_field (= 0)	stuff_bytes (0 or more)	Start of MAC Frame #1 (up to 183 bytes)	
MPEG Header (PUSI = 0)	Continuation of MAC Frame #1 (184 bytes)			
MPEG Header (PUSI = 1)	pointer_field (= M)	Tail of MAC Frame #1 (M bytes)	stuff_byte(s) (0 or more)	Start of MAC Frame #2 (M bytes)

Figure 5-6. Packet Format Where a MAC Frame Spans Multiple Packets

The Transmission Convergence sublayer must operate closely with the MAC sublayer in providing an accurate timestamp to be inserted into the Time Synchronization message (refer to Section 6.3.2 and Section 7.3).

5.6 Interaction with the Physical Layer

The MPEG-2 packet stream **MUST** be encoded according to [ITU-T J.83-B], including MPEG-2 transport framing using a parity checksum as described in [ITU-T J.83-B].

5.7 MPEG Header Synchronization and Recovery

The MPEG-2 packet stream **SHOULD** be declared “in frame” (i.e., correct packet alignment has been achieved) when five consecutive correct parity checksums, each 188 bytes from the previous one, have been received.

The MPEG-2 packet stream **SHOULD** be declared “out of frame,” and a search for correct packet alignment started, when nine consecutive incorrect parity checksums are received.

The format of MAC frames is described in detail in Section 6.

6 Media Access Control Specification

6.1 Introduction

6.1.1 Overview

This section describes version 1.1 of the DOCSIS MAC protocol. Some of the MAC protocol highlights include:

- Bandwidth allocation controlled by CMTS
- A stream of mini-slots in the upstream
- Dynamic mix of contention- and reservation-based upstream transmit opportunities
- Bandwidth efficiency through support of variable-length packets
- Extensions provided for future support of ATM or other Data PDU
- Quality-of-service including:
 - Support for Bandwidth and Latency Guarantees
 - Packet Classification
 - Dynamic Service Establishment
- Extensions provided for security at the data link layer.
- Support for a wide range of data rates.

6.1.2 Definitions

6.1.2.1 MAC-Sublayer Domain

A MAC-sublayer domain is a collection of upstream and downstream channels for which a single MAC Allocation and Management protocol operates. Its attachments include one CMTS and some number of CMs. The CMTS MUST service all of the upstream and downstream channels; each CM MAY access one or more upstream and downstream channels. The CMTS MUST police and discard any packets received that have a source MAC address that is not a unicast MAC address.¹

6.1.2.2 MAC Service Access Point

A MAC Service Access Point (MSAP) is an attachment to a MAC-sublayer domain. (Refer to Section 3.2)

6.1.2.3 Service Flows

The concept of Service Flows is central to the operation of the MAC protocol. Service Flows provide a mechanism for upstream and downstream Quality of Service management. In particular, they are integral to bandwidth allocation.

A Service Flow ID defines a particular unidirectional mapping between a CM and the CMTS. Active Upstream Service Flow IDs also have associated Service IDs or SIDs. Upstream bandwidth is allocated to SIDs, and hence to CMs, by the CMTS. Service IDs provide the mechanism by which upstream Quality of Service is implemented.

1. sentence added per rfi-n-99050 06/21/99 ew

The CMTS MAY assign one or more Service Flow IDs (SFIDs) to each CM, corresponding to the Service Flows required by the CM. This mapping can be negotiated between the CMTS and the CM during CM registration or via dynamic service establishment (refer to Section 9.4).

In a basic CM implementation, two Service Flows (one upstream, one downstream) could be used, for example, to offer best-effort IP service. However, the Service Flow concept allows for more complex CMs to be developed with support for multiple service classes while supporting interoperability with more basic modems. With these more complex modems, it is possible that certain Service Flows will be configured in such a way that they cannot carry all types of traffic. That is, they may have a maximum packet size limitation or be restricted to small fixed size unsolicited grants. Furthermore it might not be appropriate to send other kinds of data on Service Flows that are being used for Constant Bit Rate (CBR)-type applications.

Even in these complex modems, it is necessary to be able to send certain upstream packets needed for MAC management, SNMP management, key management, etc. For the network to function properly, all CMs MUST support at least one upstream and one downstream Service Flow. These Service Flows MUST always be provisioned to allow the CM to request and to send the largest possible unconcatenated MAC frame (refer to Section 6.2.2). These Service Flows are referred to as the upstream and downstream Primary Service Flows. The SID assigned to the upstream Primary Service Flow is referred to as the Primary SID.

The Primary SID MUST always be assigned to the first provisioned upstream Service Flow during the registration process (which may or may not be the same temporary SID used for the registration process). The Primary Service Flows MUST be immediately activated at registration time. The Primary SID MUST always be used for station maintenance after registration. The Primary Service Flows MAY be used for traffic. All unicast Service Flows MUST use the security association defined for the Primary Service Flow. (Refer to [DOCSIS8])

All Service Flow IDs are unique within a single MAC-sublayer domain. The length of the Service Flow ID is 32 bits. The length of the Service ID is 14 bits (although the Service ID is sometimes carried in the low-order bits of a 16-bit field).

6.1.2.4 Upstream Intervals, Mini-Slots and 6.25-Microsecond Increments

The upstream transmission time-line is divided into intervals by the upstream bandwidth allocation mechanism. Each interval is an integral number of mini-slots. A "mini-slot" is the unit of granularity for upstream transmission opportunities. There is no implication that any PDU can actually be transmitted in a single mini-slot. Each interval is labeled with a usage code which defines both the type of traffic that can be transmitted during that interval and the physical-layer modulation encoding. A mini-slot is a power-of-two multiple of 6.25µs increments, i.e., 2, 4, 8, 16, 32, 64, or 128 times 6.25µs. The relationship between mini-slots, bytes, and time ticks is described further in Section 7.3.4. The usage code values are defined in Table 6-20 and allowed use is defined in Section 6.3. The binding of these values to physical-layer parameters is defined in Table 6-18.

6.1.2.5 Frame

A frame is a unit of data exchanged between two (or more) entities at the Data Link Layer. A MAC frame consists of a MAC Header (beginning with a Frame Control byte; see Figure 6-3), and may incorporate a variable-length data PDU. The variable-length PDU includes a pair of 48-bit addresses, data, and a CRC. In special cases, the MAC Header may encapsulate multiple MAC frames (see Section 6.2.5.5) into a single MAC frame.

6.1.3 Future Use

A number of fields are defined as being “for future use” or Reserved in the various MAC frames described in this document. These fields MUST NOT be interpreted or used in any manner by this version (1.1) of the MAC protocol.

6.2 MAC Frame Formats

6.2.1 Generic MAC Frame Format

A MAC frame is the basic unit of transfer between MAC sublayers at the CMTS and the cable modem. The same basic structure is used in both the upstream and downstream directions. MAC frames are variable in length. The term “frame” is used in this context to indicate a unit of information that is passed between MAC sublayer peers. This is not to be confused with the term “framing” that indicates some fixed timing relationship.

There are three distinct regions to consider, as shown in Figure 6-1. Preceding the MAC frame is either PMD sublayer overhead (upstream) or an MPEG transmission convergence header (downstream). The first part of the MAC frame is the MAC Header. The MAC Header uniquely identifies the contents of the MAC frame. Following the header is the optional Data PDU region. The format of the Data PDU and whether it is even present is described in the MAC Header.

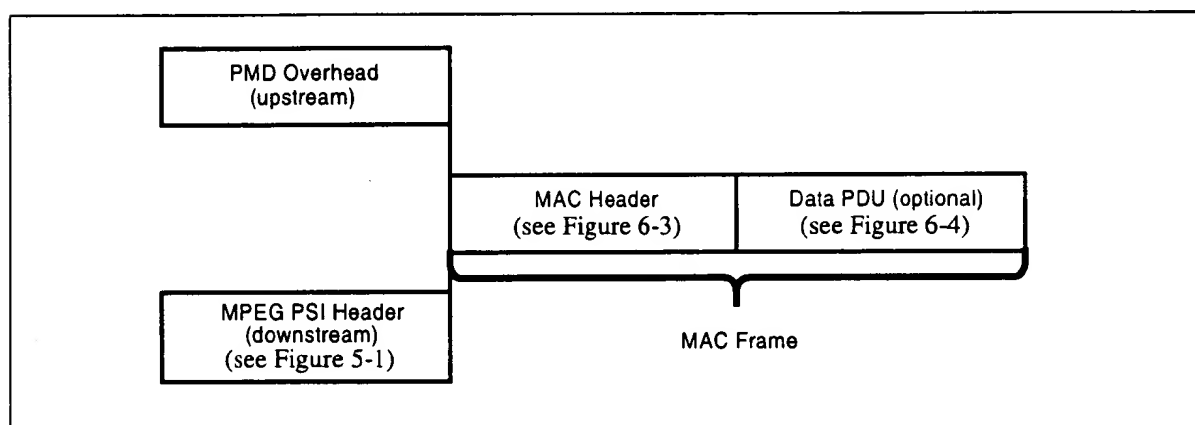


Figure 6-1. Generic MAC Frame Format

6.2.1.1 PMD Overhead

In the upstream direction, the PHY layer indicates the start of the MAC frame to the MAC sublayer. From the MAC sublayer’s perspective, it only needs to know the total amount of overhead so it can account for it in the Bandwidth Allocation process. More information on this may be found in the PMD Sublayer section of this document (Section 4).

The FEC overhead is spread throughout the MAC frame and is assumed to be transparent to the MAC data stream. The MAC sublayer does need to be able to account for the overhead when doing Bandwidth Allocation. More information on this may be found in the Upstream Bandwidth Allocation section of this document (refer to Section 7.1).

6.2.1.2 MAC Frame Transport

The transport of MAC frames by the PMD sublayer for upstream channels is shown in Figure 6-2.

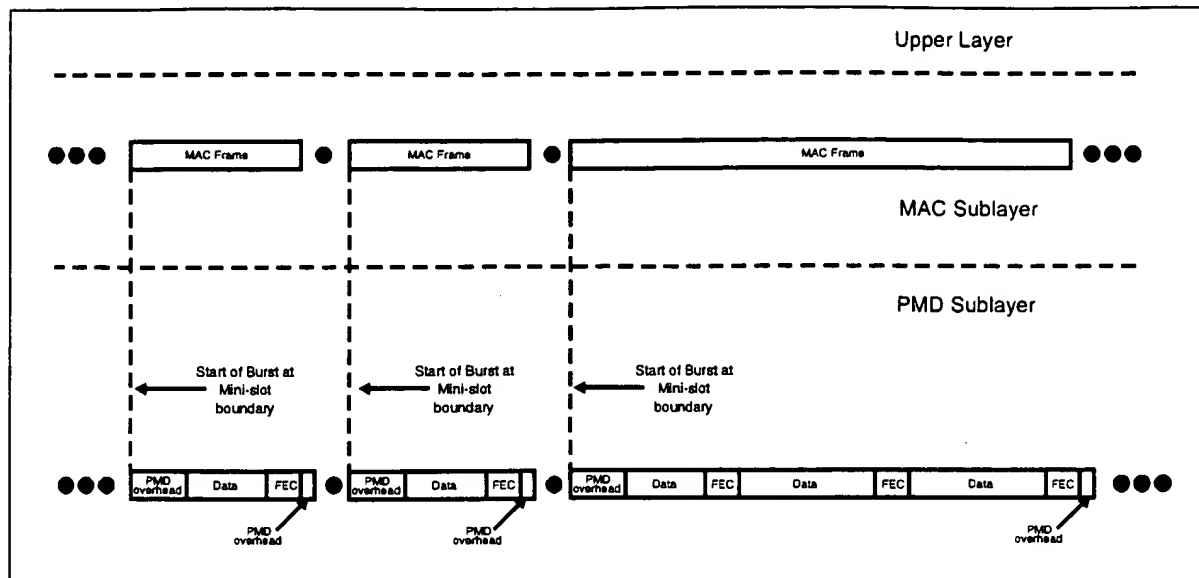


Figure 6-2. Upstream MAC/PMD Convergence

The layering of MAC frames over MPEG in the downstream channel is described in Section 5.

6.2.1.3 Ordering of Bits and Octets

Within an octet, the least-significant bit is the first transmitted on the wire. This follows the convention used by Ethernet and [ISO 8802-3]. This is often called bit-little-endian order¹.

Within the MAC layer, when numeric quantities are represented by more than one octet (i.e., 16-bit and 32-bit values), the octet containing the most-significant bits is the first transmitted on the wire. This is sometimes called byte-big-endian order.

This section follows the textual convention that when bit-fields are presented in tables, the most-significant bits are topmost in the table. For example, in Table 6-2, FC_TYPE occupies the two most-significant bits and EHDR_ON occupies the least-significant bit.

6.2.1.3.1 Representing Negative Numbers

Signed integer values will be transmitted and received in two's complement format.

6.2.1.3.2 Type-Length-Value Fields

Many MAC messages incorporate Type-Length-Value (TLV) fields. TLV fields MAY be unordered lists of TLV-tuples. Some TLV's MAY be nested (see Appendix C). All TLV Length fields MUST be greater than zero. Unless otherwise specified, Type is one byte and Length is one byte.

¹ This applies to the upstream channel only. For the downstream channel, the MPEG transmission convergence sublayer presents an octet-wide interface to the MAC, so the MAC sublayer does not define the bit order.

Using this encoding, new parameters MAY be added which some devices cannot interpret. A CM or CMTS which does not recognize a parameter type MUST skip over this parameter and MUST NOT treat the event as an error condition.

6.2.1.4 MAC Header Format

The MAC Header format MUST be as shown in Figure 6-3.

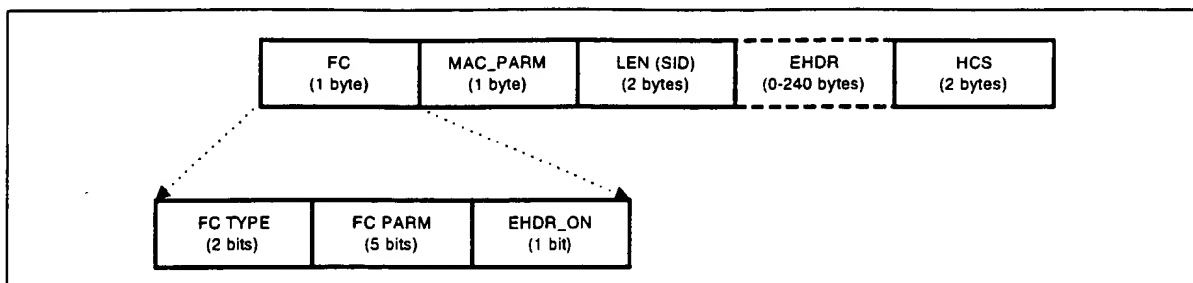


Figure 6-3. MAC Header Format

All MAC Headers MUST have the general format as shown in Table 6-1. The Frame Control (FC) field is the first byte and uniquely identifies the rest of the contents within the MAC Header. The FC field is followed by 3 bytes of MAC control; an OPTIONAL Extended Header field (EHDR); plus a Header Check Sequence (HCS) to ensure the integrity of the MAC Header.

Table 6-1. Generic MAC Header Format

MAC Header Field	Usage	Size
FC	Frame Control: Identifies type of MAC Header	8 bits
MAC_PARM	Parameter field whose use is dependent on FC: if EHDR_ON=1; used for EHDR field length (ELEN) else if for concatenated frames (see Table 6-10) used for MAC frame count else (for Requests only) indicates the number of mini-slots requested ^a	8 bits
LEN (SID)	The length of the MAC frame. The length is defined to be the sum of the number of bytes in the extended header (if present) and the number of bytes following the HCS field. (For a REQ Header, this field is the Service ID instead)	16 bits
EHDR	Extended MAC Header (where present; variable size).	0-240 bytes
HCS	MAC Header Check Sequence	2 bytes
	Length of a MAC Header	6 bytes + EHDR

a. edited 06/21/99 per rfi-n-99043

The HCS field is a 16-bit CRC that ensures the integrity of the MAC Header, even in a collision environment. The HCS field coverage MUST include the entire MAC Header, starting with the FC field and including any EHDR field that may be present. The HCS is calculated using CRC-CCITT ($x^{16} + x^{12} + x^5 + 1$) as defined in [ITU-T X.25].

The FC field is broken down into the FC_TYPE sub-field, FC_PARM sub-field and an EHDR_ON indication flag. The format of the FC field MUST be as shown in Table 6-2.

Table 6-2. FC Field Format

FC Field	Usage	Size
FC_TYPE	MAC Frame Control Type field: 00: Packet PDU MAC Header 01: ATM PDU MAC Header 10: Reserved PDU MAC Header 11: MAC Specific Header	2 bits
FC_PARM	Parameter bits, use dependent on FC_TYPE.	5 bits
EHDR_ON	When = 1, indicates that EHDR field is present. [Length of EHDR (ELEN) determined by MAC_PARM field]	1 bit

The FC_TYPE sub-field is the two MSBs of the FC field. These bits **MUST** always be interpreted in the same manner to indicate one of four possible MAC frame formats. These types include: MAC Header with Packet PDU; MAC Header with ATM cells; MAC Header reserved for future PDU types; or a MAC Header used for specific MAC control purposes. These types are spelled out in more detail in the remainder of this section.

The five bits following the FC_TYPE sub-field is the FC_PARM sub-field. The use of these bits are dependent on the type of MAC Header. The LSB of the FC field is the EHDR_ON indicator. If this bit is set, then an Extended Header (EHDR) is present. The EHDR provides a mechanism to allow the MAC Header to be extensible in an inter-operable manner.

The Transmission Convergence Sublayer stuff-byte pattern is defined to be a value of 0xFF. This precludes the use of FC byte values which have FC_TYPE = '11' and FC_PARM = '1111'.

The MAC_PARM field of the MAC Header serves several purposes depending on the FC field. If the EHDR_ON indicator is set, then the MAC_PARM field **MUST** be used as the Extended Header length (ELEN). The EHDR field **MAY** vary from 0 to 240 bytes. If this is a concatenation MAC Header, then the MAC_PARM field represents the number of MAC frames (CNT) in the concatenation (see Section 6.2.5.5). If this is a Request MAC Header (REQ) (see Section 6.2.5.3), then the MAC_PARM field represents the amount of bandwidth being requested. In all other cases, the MAC_PARM field is reserved for future use.

The third field has two possible uses. In most cases, it indicates the length (LEN) of this MAC frame. In one special case, the Request MAC Header, it is used to indicate the cable modem's Service ID since no PDU follows the MAC Header.

The Extended Header (EHDR) field provides extensions to the MAC frame format. It is used to implement data link security as well as frame fragmentation, and can be extended to add support for additional functions in future releases. Initial implementations **SHOULD** pass this field to the processor. This will allow future software upgrades to take advantage of this capability. (Refer to Section 6.2.6, "Extended MAC Headers" for details.)

6.2.1.5 Data PDU

The MAC Header **MAY** be followed by a Data PDU. The type and format of the Data PDU is defined in the Frame Control field of the MAC Header. The FC field explicitly defines a Packet Data PDU, an ATM Data PDU, a MAC-Specific Frame and a reserved code point (used as an escape mechanism for future extensions). All CMs **MUST** use the length in the MAC Header to skip over any reserved data.

6.2.2 Packet-Based MAC Frames

6.2.2.1 Variable-Length Packets

The MAC sublayer **MUST** support a variable-length Ethernet/[ISO8802-3]-type Packet Data PDU. Normally, the Packet PDU **MUST** be passed across the network in its entirety, including its original CRC.¹ A unique Packet MAC Header is appended to the beginning. The frame format without an Extended header **MUST** be as shown in Figure 6-4 and Table 6-3.

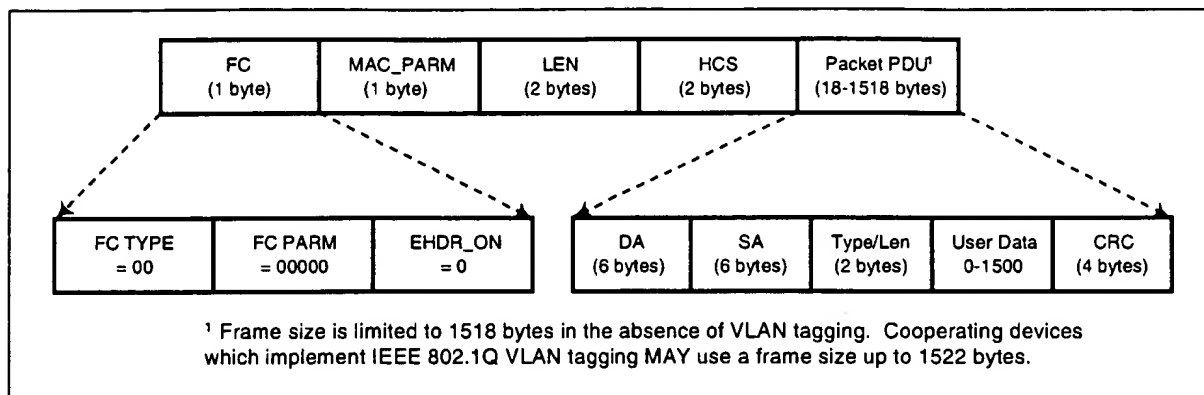


Figure 6-4. Ethernet/802.3 Packet PDU Format

1. The one exception is the case of Payload Header Suppression. In this case, all bytes except those suppressed **MUST** be passed across the network and the CRC covers only those bytes actually transmitted. (Refer to Section 6.2.6.3.1)

Table 6-3. Example Packet PDU Format

Field	Usage	Size
FC	FC_TYPE = 00; Packet MAC Header FC_PARM[4:0] = 00000; other values reserved for future use and ignored EHDR_ON = 0; No EHDR present this example	8 bits
MAC_PARM	Reserved, MUST be set to zero if there is no EHDR; Otherwise set to length of EHDR	8 bits
LEN	LEN = n; length of Packet PDU in bytes	16 bits
EHDR	EHDR = 0; Extended MAC Header not present in this example	0 bytes
HCS	MAC Header Check Sequence	2 bytes
Packet Data	Packet PDU: DA - 48 bit Destination Address SA - 48 bit Source Address Type/Len - 16 bit Ethernet Type or [ISO8802-3] Length Field User Data (variable length, 0-1500 bytes) CRC - 32-bit CRC over packet PDU (as defined in Ethernet/[ISO8802-3])	n bytes
	Length of example Packet MAC frame	6 + n bytes

Under certain circumstances (see appendix M) it may be necessary to transmit a packet PDU MAC frame without an actual PDU. This is done so that the extended header can be used to carry certain information about the state of the service flow. Such a frame will have the length field in MAC header set to 0 and will have no packet data, which means no CRC.¹

1. paragraph added 6.8.99 per rfi-n-99039

6.2.3 ATM Cell MAC Frames

The FC_TYPE 0x01 is reserved for future definition of ATM Cell MAC Frames. This FC_TYPE field in the MAC Header indicates that an ATM PDU is present. This PDU MUST be silently discarded by MAC implementations of this version (1.1) of the specification. Compliant version 1.1 implementations MUST use the length field to skip over the ATM PDU.¹

6.2.4 Reserved PDU MAC Frames

The MAC sublayer provides a reserved FC code point to allow for support of future (to be defined) PDU formats. The FC field of the MAC Header indicates that a Reserved PDU is present. This PDU MUST be silently discarded by MAC implementations of this version (1.1) of the specification. Compliant version 1.1 implementations MUST use the length field to skip over the Reserved PDU.

The format of the Reserved PDU without an extended header MUST be as shown in Figure 6-5 and Table 6-4.

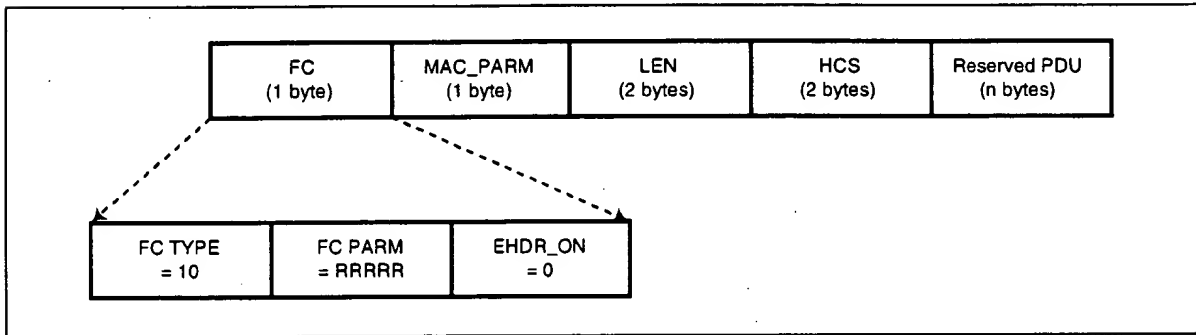


Figure 6-5. Reserved PDU Format

Table 6-4. Example Reserved PDU Format

Field	Usage	Size
FC	FC_TYPE = 10; Reserved PDU MAC Header FC_PARM[4:0]; reserved for future use EHDR_ON = 0; No EHDR present this example	8 bits
MAC_PARM	Reserved for future use	8 bits
LEN	LEN = n; length of Reserved PDU in bytes	16 bits
EHDR	EHDR = 0; Extended MAC Header not present this example	0 bytes
HCS	MAC Header Check Sequence	2 bytes
User Data	Reserved Data PDU	n bytes
Length of a Reserved PDU MAC frame		6 + n bytes

1. edited 06/21/99 per rfi-n-99043 ew

6.2.5 MAC-Specific Headers

There are several MAC Headers which are used for very specific functions. These functions include support for downstream timing and upstream ranging/power adjust, requesting bandwidth, fragmentation and concatenating multiple MAC frames.

Table 6-5 describes FC_PARM usage within the MAC Specific Header.

FC_PARM	Header/Frame Type
00000	Timing Header
00001	MAC Management Header
00010	Request Frame
00011	Fragmentation Header
11100	Concatenation Header

Table 6-5. MAC-Specific Headers and Frames

6.2.5.1 Timing Header

A specific MAC Header is identified to help support the timing and adjustments required. In the downstream, this MAC Header **MUST** be used to transport the Global Timing Reference to which all cable modems synchronize. In the upstream, this MAC Header **MUST** be used as part of the Ranging message needed for a cable modem's timing and power adjustments. The Timing MAC Header is followed by a Packet Data PDU. The format **MUST** be as shown in Figure 6-6 and Table 6-6.

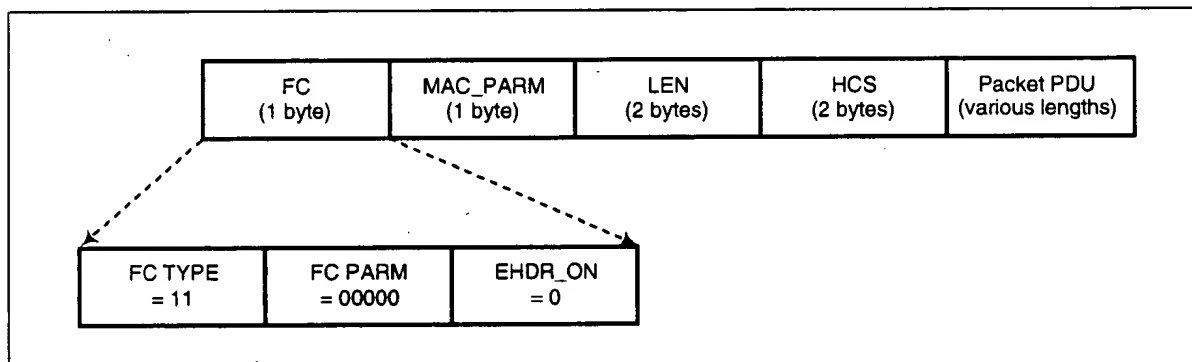


Figure 6-6. Timing MAC Header

Table 6-6. Timing MAC Header Format

Field	Usage	Size
FC	FC_TYPE = 11; MAC Specific Header FC_PARM[4:0] = 00000; Timing MAC Header EHDR_ON = 0; Extended header prohibited for SYNC and RNG-REQ	8 bits
MAC_PARM	Reserved for future use	8 bits
LEN	LEN = n; Length of Packet PDU in bytes	16 bits
EHDR	Extended MAC Header not present	0 bytes
HCS	MAC Header Check Sequence	2 bytes
Packet Data	MAC Management message: SYNC message (downstream only) RNG-REQ (upstream only)	n bytes
	Length of Timing Message MAC frame	6 + n bytes

6.2.5.2 MAC Management Header

A specific MAC Header is identified to help support the MAC management messages required. This MAC Header MUST be used to transport all MAC management messages (refer to Section 6.3). The format MUST be as shown Figure 6-7 and Table 6-7.

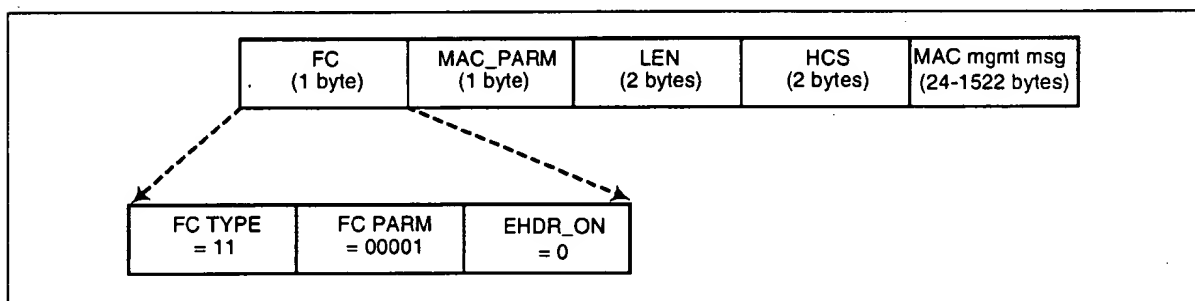
Figure 6-7. Management MAC Header¹

Table 6-7. Example Management MAC Header Format

Field	Usage	Size
FC	FC_TYPE = 11; MAC Specific Header FC_PARM[4:0] = 00001; Management MAC Header EHDR_ON=0; No EHDR present this example	8 bits
MAC_PARM	Reserved for future use	8 bits
LEN	LEN = n; length of Packet PDU in bytes	16 bits
EHDR	Extended MAC Header not present this example	0 bytes
HCS	MAC Header Check Sequence	2 bytes
Packet Data	MAC Management message:	n bytes
	Length of Example Management MAC frame	6 + n bytes

1. Figure 6-7 edited 6/7/99 per ECN rfi-n-99035

6.2.5.3 Request Frame

The Request Frame is the basic mechanism that a cable modem uses to request bandwidth. As such, it is only applicable in the upstream. There **MUST** be no Data PDUs following the Request Frame. The general format of the Request **MUST** be as shown in Figure 6-8 and Table 6-8.




Figure 6-8. Request Frame Format

Table 6-8. Request Frame (REQ) Format

Field	Usage	Size
FC	FC_TYPE = 11; MAC-Specific Header FC_PARM[4:0] = 00010; MAC Header only; no data PDU following EHDR_ON = 0; No EHDR allowed	8 bits
MAC_PARM	REQ, total number of minislots requested	8 bits
SID	Service ID (0...0x1FFF) ^a	16 bits
EHDR	Extended MAC Header not allowed	0 bytes
HCS	MAC Header Check Sequence	2 bytes
	Length of a REQ MAC Header	6 bytes

a. edited per rfi-n-99043 06/21/99 ew

Because the Request Frame does not have a Data PDU following it, the LEN field is not needed. The LEN field **MUST** be replaced with an SID. The SID **MUST** uniquely identify a particular Service Flow within a given CM.

The bandwidth request, REQ, **MUST** be specified in mini-slots. The REQ field **MUST** indicate the current total amount of bandwidth requested for this service queue including appropriate allowance for the PHY overhead.

6.2.5.4 Fragmentation Header

The Fragmentation MAC Header provides the basic mechanism to split a larger MAC PDU into smaller pieces that are transmitted individually and then re-assembled at the CMTS. As such, it is only applicable in the upstream. The general format of the Fragmentation MAC Header **MUST** be as shown in Figure 6-9.

A compliant CM **MUST** support fragmentation. A compliant CMTS **MAY** support fragmentation. To decrease the burden on the CMTS and to reduce unnecessary overhead, fragmentation headers **MUST NOT** be used on unfragmented frames.

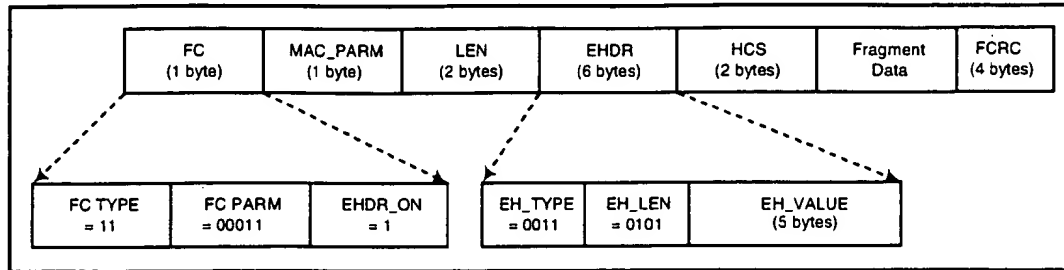


Figure 6-9. Fragmentation MAC Header Format

Table 6-9. Fragmentation MAC Frame (FRAG) Format

Field	Usage	Size
FC	FC_TYPE = 11; MAC-Specific Header FC_PARM [4:0] = 00011; Fragmentation MAC Header EHDR_ON = 1; Fragmentation EHDR follows	8 bits
MAC_PARM	ELEN = 6 bytes; length of Fragmentation EHDR	8 bits
LEN	LEN = length of fragment payload + EHDR length + FCRC length	16 bits
EHDR	Refer to Section 6.2.6.2	6 bytes
HCS	MAC Header Check Sequence	2 bytes
Fragment Data	Fragment payload; portion of total MAC PDU being sent	n bytes
FCRC	CRC - 32-bit CRC over Fragment Data payload (as defined in Ethernet/[ISO8802-3])	4 bytes
	Length of a MAC Fragment Frame	16 + n bytes

6.2.5.5 Concatenation Header

A Specific MAC Header is defined to allow multiple MAC frames to be concatenated. This allows a single MAC “burst” to be transferred across the network. The PHY overhead¹ and the Concatenation MAC Header only occur once. Concatenation of multiple MAC frames **MUST** be as shown in Figure 6-10.

A compliant CM **MUST** support concatenation. A compliant CMTS **MAY** support concatenation. Concatenation only applies to upstream traffic. Concatenation **MUST NOT** be used on downstream traffic.

1. This includes the preamble, guard time, and possibly zero-fill bytes in the last codeword. The FEC overhead recurs for each codeword.




Figure 6-10. Concatenation of Multiple MAC Frames

Only one Concatenation MAC Header **MUST** be present per MAC “burst.” Nested concatenation **MUST NOT** be allowed. Immediately following the Concatenation MAC Header **MUST** be the MAC Header of the first MAC frame. Information within the MAC Header indicates the length of the first MAC Frame and provides a means to find the start of the next MAC Frame. Each MAC frame within a concatenation **MUST** be unique and **MAY** be of any type. This means that Packet and MAC-specific Frames **MAY** be mixed together. However, all frames in a concatenation **MUST** be assigned to the same Service Flow.

The embedded MAC frames **MAY** be addressed to different destinations and **MUST** be delivered as if they were transmitted individually.

The format of the Concatenation MAC Header **MUST** be as shown in Figure 6-11 and Table 6-10.

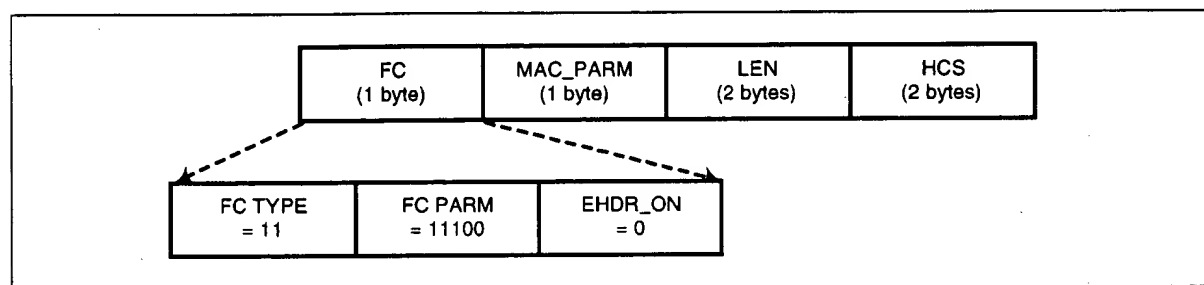


Figure 6-11. Concatenation MAC Header Format

Table 6-10. Concatenated MAC Frame Format

Field	Usage	Size
FC	FC_TYPE = 11; MAC Specific Header FC_PARM[4:0] = 11100; Concatenation MAC Header EHDR_ON = 0; No EHDR with Concatenation Header	8 bits
MAC_PARM	CNT, number of MAC frames in this concatenation CNT = 0 indicates unspecified number of MAC frames	8 bits
LEN	LEN = x + ... + y; length of all following MAC frames in bytes	16 bits
EHDR	Extended MAC Header MUST NOT be used	0 bytes
HCS	MAC Header Check Sequence	2 bytes
MAC frame 1	First MAC frame: MAC Header plus OPTIONAL data PDU	x bytes
MAC frame n	Last MAC frame: MAC Header plus OPTIONAL data PDU	y bytes
	Length of Concatenated MAC frame	6 + LEN bytes

The MAC_PARM field in the Concatenation MAC header provides a count of MAC frames as opposed to EHDR length or REQ amount as used in other MAC headers. If the field is non-zero, then it **MUST** indicate the total count of MAC Frames (CNT) in this concatenation burst.

6.2.6 Extended MAC Headers

Every MAC Header, except the Timing, Concatenation MAC Header and Request Frame, has the capability of defining an Extended Header field (EHDR). The presence of an EHDR field **MUST** be indicated by the EHDR_ON flag in the FC field being set. Whenever this bit is set, then the MAC_PARM field **MUST** be used as the EHDR length (ELEN). The minimum defined EHDR is 1 byte. The maximum EHDR length is 240 bytes.

A compliant CMTS & CM **MUST** support extended headers.

The format of a generic MAC Header with an Extended Header included **MUST** be as shown in Figure 6-12 and Table 6-11. Note: Extended Headers **MUST NOT** be used in a Concatenation MAC Header, but **MAY** be included as part of the MAC Headers within the concatenation.

Extended Headers **MUST NOT** be used in Request Frames and Timing MAC Headers.

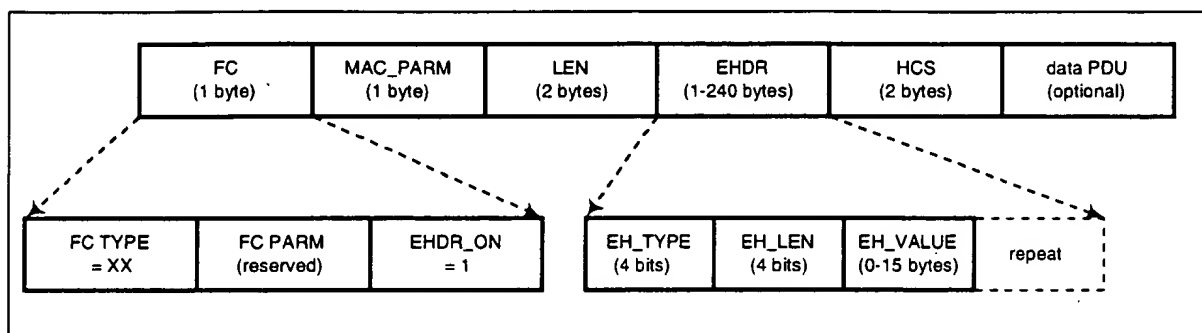


Figure 6-12. Extended MAC Format

Table 6-11. Example Extended Header Format

Field	Usage	Size
FC	FC_TYPE = XX; Applies to all MAC Headers FC_PARM[4:0] = XXXXX; dependent on FC_TYPE EHDR_ON = 1; EHDR present this example	8 bits
MAC_PARM	ELEN = x; length of EHDR in bytes	8 bits
LEN	LEN = x + y; length of EHDR plus OPTIONAL data PDU in bytes	16 bits
EHDR	Extended MAC Header present this example	x bytes
HCS	MAC Header Check Sequence	2 bytes
PDU	OPTIONAL data PDU	y bytes
	Length of MAC frame with EHDR	6 + x + y bytes

Since the EHDR increases the length of the MAC frame, the LEN field **MUST** be increased to include both the length of the Data PDU and the length of the EHDR.

The EHDR field consists of one or more EH elements. Each EH element is variable sized. The first byte of the EH element **MUST** contain a type and a length field. Every CM **MUST** use this length to skip over any unknown EH elements. The format of an EH element **MUST** be as shown in Table 6-12.

Table 6-12. EH Element Format

EH Element Fields	Usage	Size
EH_TYPE	EH element Type Field	4 bits
EH_LEN	Length of EH_VALUE	4 bits
EH_VALUE	EH element data	0-15 bytes

The types of EH element defined in Table 6-13 **MUST** be supported. Reserved and extended types are undefined at this point and **MUST** be ignored.

The first ten EH element types are intended for one-way transfer between the cable modem and the CMTS. The next five EH element types are for end-to-end usage within a MAC-sublayer domain. Thus, the information attached to the EHDR on the upstream **MUST** also be attached when the information is forwarded. The final EH element type is an escape mechanism that allows for more types and longer values, and **MUST** be as shown in Table 6-13.

Table 6-13. Extended Header Types

EH_TYPE	EH_LEN	EH_VALUE
0	0	Null configuration setting; may be used to pad the extended header. The EH_LEN MUST be zero, but the configuration setting may be repeated.
1	3	Request: mini-slots requested (1 byte); SID (2 bytes) [CM --> CMTS]
2	2	Acknowledgment requested; SID (2 bytes) [CM --> CMTS]
3 (= BP_UP)	4	Upstream Privacy EH Element [DOCSIS8]
	5	Upstream Privacy with Fragmentation ^a EH Element [DOCSIS8](See 6.2.7)
4 (= BP_DOWN)	4	Downstream Privacy EH Element [DOCSIS8]
5	1	Service Flow EH Element; Payload Header Suppression Header
	2	Service Flow EH Element; Payload Header Suppression Header (1 byte) Unsolicited Grant Synchronization Header (1 byte)
6 - 9		Reserved
10 - 14		Reserved [CM <-> CM]
15	XX	Extended EH Element: EHX_TYPE (1 byte), EHX_LEN (1 byte), EH_VALUE (length determined by EHX_LEN)

a. An Upstream Privacy with Fragmentation EH Element **MUST** only occur within a Fragmentation MAC-Specific Header. (Refer to Section 6.2.5.4)

6.2.6.1 Piggyback Requests

Several Extended Headers can be used to request bandwidth for subsequent transmissions. These requests are generically referred to as "piggyback requests". They are extremely valuable for performance because they are not subject to contention as Request Frames generally are. (Refer to Section 7.4)

Requests for additional bandwidth can be included in Request, Upstream Privacy and Upstream Privacy with Fragmentation Extended Header elements.

6.2.6.2 Fragmentation Extended Header

Fragmented packets use a combination of the Fragmentation MAC header and a modified version of the Upstream Privacy Extended header. Section 6.2.5.4 describes the Fragmentation MAC header. The Upstream Privacy Extended Header with Fragmentation, also known as the Fragmentation Extended Header, **MUST** be as shown in Table 6-14.

Table 6-14. Fragmentation Extended Header Format

EH Element Fields	Usage	Size
EH_TYPE	Upstream Privacy EH element = 3	4 bits
EH_LEN	Length of EH_VALUE = 5	4 bits
EH_VALUE	Key_seq; same as in BP_UP	4 bits
	Ver = 1; version number for this EHDR	4 bits
	BPI_ENABLE If BPI_ENABLE=0, BPI disabled If BPI_ENABLE=1, BPI enabled	1 bit
	Toggle bit; same as in BP_UP ^a	1 bit
	SID; Service ID associated with this fragment	14 bits
	REQ; number of mini-slots for a piggyback request	8 bits
	Reserved; must be set to zero	2 bits
	First_Frag; set to one for first fragment only	1 bit
	Last_Frag; set to one for last fragment only	1 bit
	Frag_seq; fragment sequence count, incremented for each fragment.	4 bits

a. Refer to [DOCSIS8]

6.2.6.3 Service Flow Extended Header

The Service Flow EH Element is used to enhance Service Flow operations. It may consist of one or two bytes in the EH_VALUE field. The Payload Header Suppression Header is the only byte in a one byte field or the first byte in a two byte field. The Unsolicited Grant Synchronization Header is the second byte in a two byte field.

6.2.6.3.1 Payload Header Suppression Header

In Payload Header Suppression (PHS), a repetitive portion of the payload headers following the HCS is suppressed by the sending entity and restored by the receiving entity. In the upstream, the sending entity is the CM and the receiving entity is the CMTS. In the downstream, the sending entity is the CMTS and the receiving entity is the CM.

For small payloads, Payload Header Suppression provides increased bandwidth efficiency without having to use compression. Payload Header Suppression may be separately provisioned in the upstream and downstream, and is referenced with an extended header element.

A compliant CM **MUST** support Payload Header Suppression.¹ A compliant CMTS **MAY** support Payload Header Suppression.

1. This is not intended to imply that the CM must be capable of determining when to invoke Payload Header Suppression. Payload Header Suppression support is only required for the explicitly signalled case.

The Payload Header Suppression Extended Header sub-element has the following format:

Table 6-15. Payload Header Suppression EHDR Sub-Element Format

EH Element Fields	Usage		Size
EH_TYPE	Service Flow EH_TYPE = 5		4 bits
EH_LEN	Length of EH_VALUE = 1		4 bits
EH_VALUE	0	Indicates no payload header suppression on current packet.	8 bits
	1-255	Payload Header Suppression Index (PHSI)	

The Payload Header Suppression Index is unique per SID in the upstream and unique per CM in the downstream. Payload Header Suppression is disabled if this Extended Header element is omitted or, if included, with the PHSI value set to 0. The Payload Header Suppression Index (PHSI) references the suppressed byte string known as a Payload Header Suppression Field (PHSF).

Note: While PHS signaling allows for up to 254 Payload Header Suppression Rules per Service Flow, the exact number of PHS rules supported per Service Flow is implementation dependent. Similarly, PHS signaling allows for PHS Sizes of up to 255 bytes, however, the maximum PHS Size supported is implementation dependent. For interoperability, the minimum PHS Size that MUST be supported is 64 bytes for any PHS rule supported. As with any other parameter requested in a Dynamic Service Request, a PHS-related DSx request can be denied because of a lack of resources.¹

The Upstream Suppression Field MUST begin with the first byte following the MAC Header Checksum. The Downstream Suppression Field MUST begin with the thirteenth byte following the MAC Header Checksum. This allows the Ethernet SA and DA to be available for filtering by the CM.

The operation of Baseline Privacy (refer to [DOCSIS8]) is not affected by the use of PHS. When Fragmentation is inactive, Baseline Privacy begins encryption and decryption with the thirteenth byte following the MAC Header checksum.

The Packet PDU CRC is always transmitted, and MUST be calculated only on the bytes transmitted. The bytes that are suppressed MUST NOT be included in the CRC calculation.

6.2.6.3.2 Unsolicited Grant Synchronization Header

The Unsolicited Grant Synchronization Header may be used to pass status information regarding Service Flow scheduling between the CM and CMTS. It is currently only defined for use in the upstream with Unsolicited Grant and Unsolicited Grant with Activity Detection scheduling services. (Refer to Section 8.2).

This extended header is similar to the Payload Suppression EHDR except that the EH_LEN is 2, and the EH_VALUE has one additional byte which includes information related to Unsolicited Grant Synchronization. For all other Service Flow Scheduling Types, the field SHOULD NOT be included in the Extended Header Element generated by the CM. The CMTS MAY ignore this field.

1. Footnote edited 07/01/99 per rfi-n-99019. ew

Table 6-16. Unsolicited Grant Synchronization EHDR Sub-Element Format

EH Element Fields	Usage		Size
EH_TYPE	Service Flow EH_TYPE = 5		4 bits
EH_LEN	Length of EH_VALUE = 2		4 bits
EH_VALUE	0	Indicates no payload header suppression on current packet.	8 bits [always present]
	1-255	Payload Header Suppression Index (PHSI)	
	Queue Indicator		1 bit
	Active Grants		7 bits

6.2.7 Fragmented MAC Frames

When enabled, fragmentation is initiated any time the grant length is less than the requested length. This normally occurs because the CMTS chooses to grant less than the requested bandwidth.

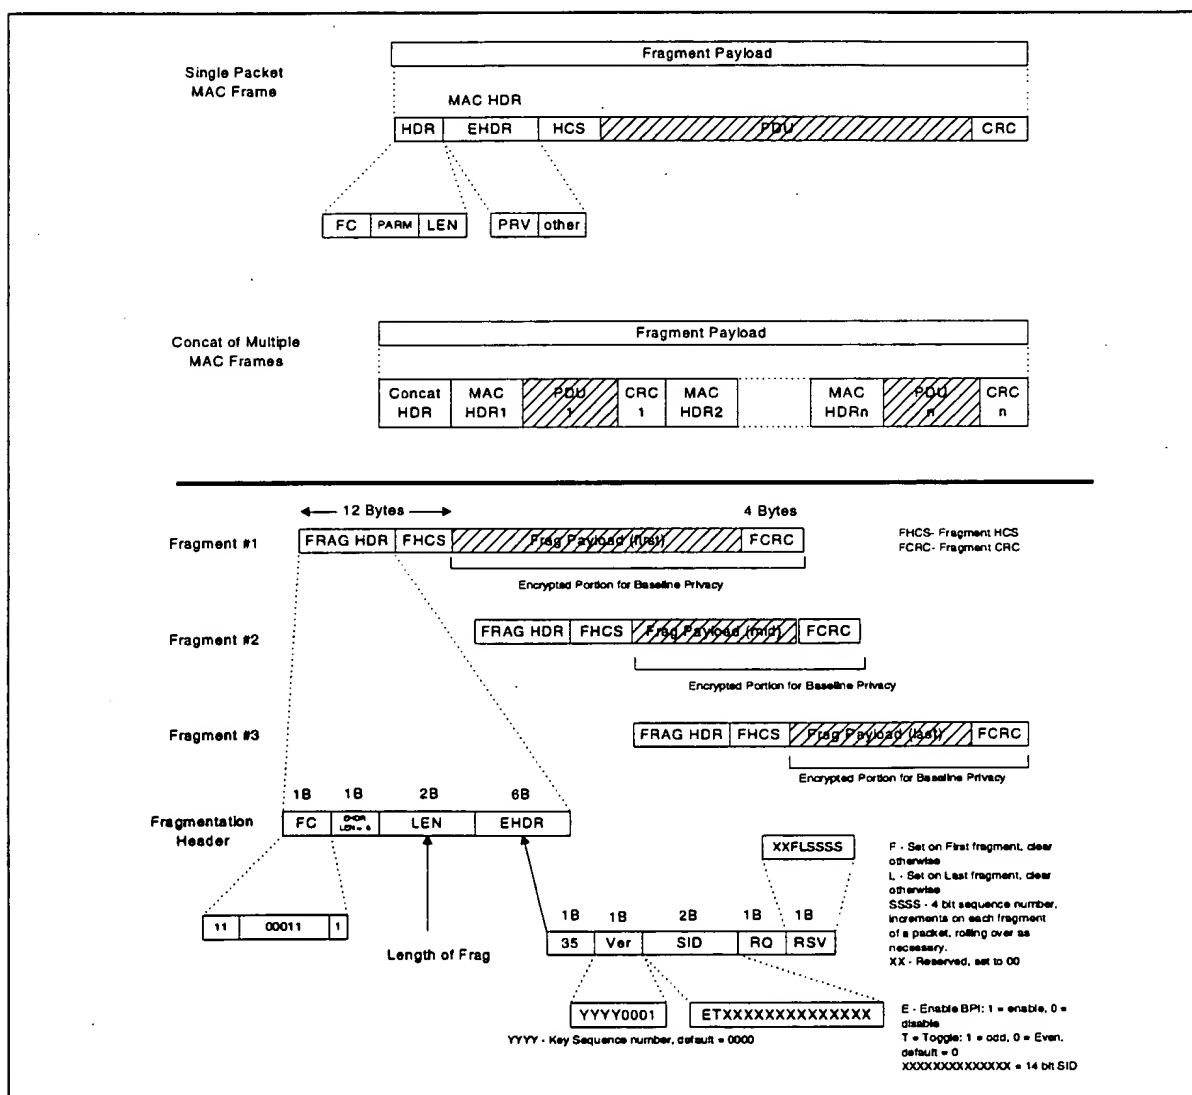


Figure 6-13. Fragmentation Details

The CM MAC calculates how many bytes of the original frame, including overhead for a fragmentation header and CRC, can be sent in the received grant. The CM MAC generates a fragmentation header for each fragment. Fragmented frames use the MAC Message type (FC = 11). The FC parameter field is set to (00011), in order to uniquely identify the fragmentation header from other MAC Message types. A four bit sequence field is used in the last byte of the Extended Header field to aid in reassembly and to detect dropped or missing fragments. The CM arbitrarily selects a sequence number for the first fragment of a frame.¹ Once the sequence number is selected for the first fragment, the CM increments the sequence number by one for each fragment transmitted for that frame. There are two flags associated with the sequence number, F and L, where F is set to indicate the first fragment and L is set to indicate the last fragment. Both are cleared for middle fragments. The CMTS stores the sequence number of the first fragment (F bit set) of each frame. The CMTS MUST verify that the fragment sequence field increments (by one) for each fragment of the frame.

The REQ field in the fragmentation header is used by the fragmentation protocol for First and Middle fragments (refer to Section 8.3). For the Last fragment, the REQ field is interpreted as a request for bandwidth for a subsequent frame.

Fragmentation headers are fixed size and MUST contain only a Fragmentation extended header element. The extended header consists of a Privacy EH element extended by one byte to make the fragment overhead an even 16 bytes. A Privacy EH element is used whether the original packet header contained a Privacy EH element or not. If privacy is in use, Key Sequence number, Version, Enable bit, Toggle bit and SID in the fragment EH element are the same with those of BP EH element inside the original MAC frame. If privacy is not in use, the Privacy EH element is used but the enable bit is cleared. The SID used in the fragment EH element MUST match the SID used in the Partial Grant that initiated the fragmentation. The same extended header must be used for all fragments of a packet. A separate CRC must be calculated for each fragment (note that each MAC frame payload will also contain the CRC for that packet). A packet CRC of a reassembled packet MAY be checked by the CMTS even though an FCRC covers each fragment.

The CMTS MUST make certain that any fragmentary grant it makes is large enough to hold at least 17 bytes of MAC layer data. This is to ensure that the grant is large enough to accommodate fragmentation overhead plus at least 1 byte of actual data. The CMTS may want to enforce an even higher limit as small fragments are extremely inefficient.

When Fragmentation is active, Baseline Privacy encryption and decryption begin with the first byte following the MAC Header checksum.

6.2.7.1 Considerations for Concatenated Packets and Fragmentation

MAC Management Messages and Data PDUs can occur in the same concatenated frame. Without fragmentation, the MAC Management Messages within a concatenated frame would be unencrypted. However, with fragmentation enabled on the concatenated frame, the entire concatenated frame is encrypted based on the Privacy Extended Header Element. This allows Baseline Privacy to encrypt each fragment without examining its contents. Clearly, this only applies when Baseline Privacy is enabled.

To ensure encryption synchronization, if fragmentation, concatenation and Baseline Privacy are all enabled, a CM MUST NOT concatenate BPKM MAC Management messages. This ensures that BPKM MAC Management messages are always sent unencrypted.

1. Note, 'frame' always refers to either frames with a single Packet PDU or concatenated frame.

6.2.8 Error-Handling

The cable network is a potentially harsh environment that can cause several different error conditions to occur. This section, together with Section 9.5, describes the procedures that are required when an exception occurs at the MAC framing level.

The most obvious type of error occurs when the HCS on the MAC Header fails. This can be a result of noise on the network or possibly by collisions in the upstream channel. Framing recovery on the downstream channel is performed by the MPEG transmission convergence sublayer. In the upstream channel, framing is recovered on each transmitted burst, such that framing on one burst is independent of framing on prior bursts. Hence, framing errors within a burst are handled by simply ignoring that burst; i.e., errors are unrecoverable until the next burst.

A second exception, which applies only to the upstream, occurs when the Length field is corrupted and the MAC thinks the frame is longer or shorter than it actually is. Synchronization will recover at the next valid upstream data interval.

For every MAC transmission, The HCS **MUST** be verified. When a bad HCS is detected, the MAC Header and any payload **MUST** be dropped.

For Packet PDU transmissions, a bad CRC **MAY** be detected. Since the CRC only covers the Data PDU and the HCS covers the MAC Header, the MAC Header is still considered valid. Thus, the Packet PDU **MUST** be dropped, but any pertinent information in the MAC Header (e.g., bandwidth request information) **MAY** be used.

6.2.8.1 Error Recovery During Fragmentation

There are some special error handling considerations for fragmentation. Each fragment has its own fragmentation header complete with an HCS and its own FCRC. There may be other MAC headers and CRCs within the fragmented payload. However, only the HCS of the fragment header and the FCRC are used for error detection during fragment reassembly.

If the HCS for a fragment fails the CMTS **MUST** discard that fragment. If the HCS passes but the FCRC fails, the CMTS **MUST** discard that fragment, but **MAY** process any requests in the fragment header. The CMTS **SHOULD** process such a request if it is performing fragmentation in Piggyback Mode. (Refer to Section 8.3.2.2) This allows the remainder of the frame to be transmitted as quickly as possible.

If a CMTS is performing fragmentation in Multiple Grant Mode (refer to Section 8.3.2.1) it **SHOULD** complete all the grants necessary to fulfil the CM's original request even if a fragment is lost or discarded. This allows the remainder of the frame to be transmitted as quickly as possible.

If any fragment of a non-concatenated MAC frame is lost or discarded the CMTS **MUST** discard the rest of that frame. If a fragment of a concatenated MAC frame is lost or discarded the CMTS **MAY** forward any frames within the concatenation that have been received correctly or it **MAY** discard all the frames in the concatenation.

A CMTS **MUST** terminate fragment reassembly if any of the following occurs for any fragment on a given SID:

- The CMTS receives a fragment with the L bit set.
- The CMTS receives an upstream fragment, other than the first one, with the F bit set.
- The CMTS receives a packet PDU frame with no fragmentation header.¹
- The CMTS deletes the SID for any reason.

1. edited 06/21/99 per rfi-n-99043

In addition, the CMTS MAY terminate fragment reassembly based on implementation dependent criteria such as a reassembly timer. When a CMTS terminates fragment reassembly it MUST dispose of (either by discarding or forwarding) the reassembled frame(s).

6.2.8.2 Error Codes and Messages

Appendix J lists CM and CMTS error codes and messages. When reporting error conditions, these codes MUST be used as indicated in [DOCSIS5] and MAY be used for reporting errors via vendor-specific interfaces. If the error codes are used, the error messages MAY be replaced by other descriptive messages.

6.3 MAC Management Messages

6.3.1 MAC Management Message Header

MAC Management Messages MUST be encapsulated in an LLC unnumbered information frame per [ISO8802-2], which in turn is encapsulated within the cable network MAC framing, as shown in Figure 6-14. Figure 6-14 shows the MAC Header and the MAC Management Message Header fields which are common across all MAC Management Messages.

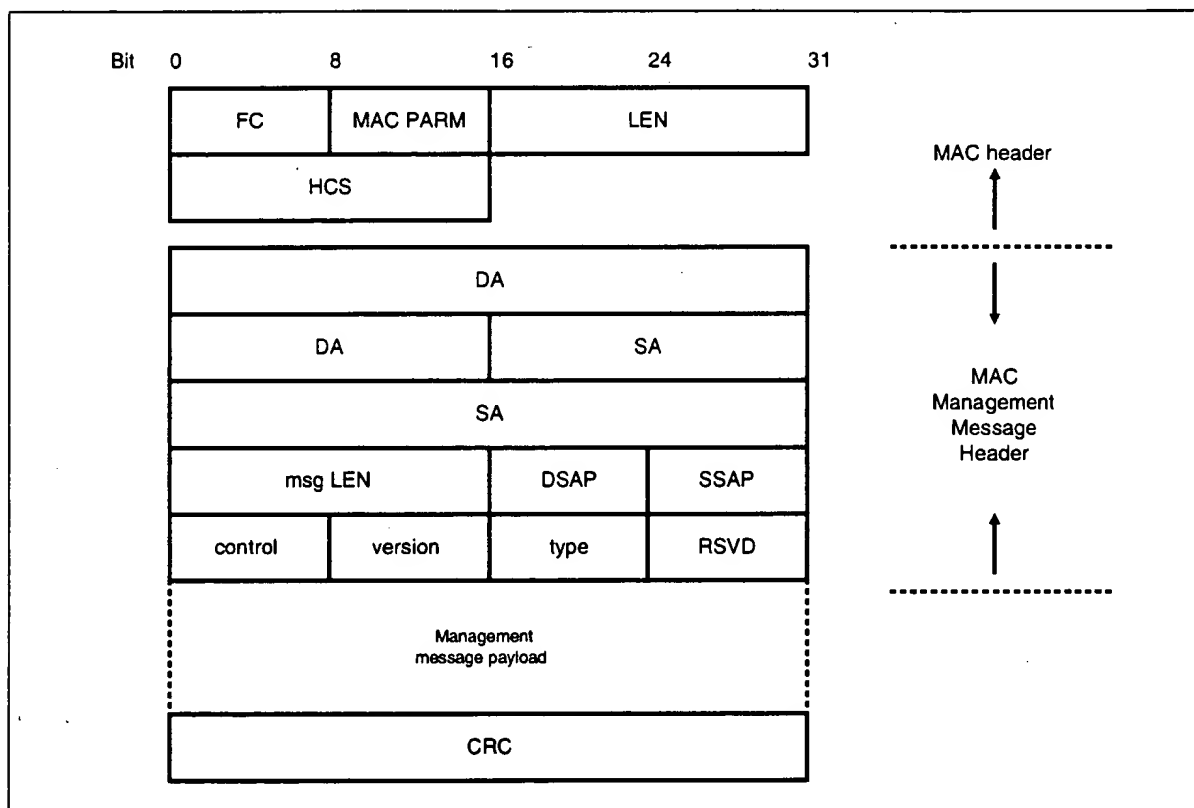


Figure 6-14. MAC Header and MAC Management Message Header Fields

The fields MUST be as defined below.

FC, MAC PARM, LEN, HCS Common MAC frame header -refer to Section 6.2.1.4 for details. All messages use a MAC-specific header.

Destination Address (DA)	MAC management frames will be addressed to a specific CM unicast address or to the DOCSIS management multicast address. These DOCSIS MAC management addresses are described in Appendix A.
Source Address (SA)	The MAC address of the source CM or CMTS system.
Msg Length	Length of the MAC message from DSAP to the end of the payload.
DSAP	The LLC null destination SAP (00) as defined by [ISO8802-2].
SSAP	The LLC null source SAP (00) as defined by [ISO8802-2].
Control	Unnumbered information frame (03) as defined by [ISO8802-2].
Version & Type	Each 1 octet. Refer to Table 6-17.

Table 6-17. MAC Management Message Types

Type Value	Version	Message Name	Message Description
1	1	SYNC	Timing Synchronization
2	1	UCD	Upstream Channel Descriptor
3	1	MAP	Upstream Bandwidth Allocation
4	1	RNG-REQ	Ranging Request
5	1	RNG-RSP	Ranging Response
6	1	REG-REQ	Registration Request
7	1	REG-RSP	Registration Response
8	1	UCC-REQ	Upstream Channel Change Request
9	1	UCC-RSP	Upstream Channel Change Response
10	1	TRI-TCD	Telephony Channel Descriptor [DOCSIS6]
11	1	TRI-TSI	Termination System Information [DOCSIS6]
12	1	BPKM-REQ	Privacy Key Management Request [DOCSIS8]
13	1	BPKM-RSP	Privacy Key Management Response [DOCSIS8]
14	2	REG-ACK	Registration Acknowledge
15	2	DSA-REQ	Dynamic Service Addition Request
16	2	DSA-RSP	Dynamic Service Addition Response
17	2	DSA-ACK	Dynamic Service Addition Acknowledge
18	2	DSC-REQ	Dynamic Service Change Request
19	2	DSC-RSP	Dynamic Service Change Response
20	2	DSC-ACK	Dynamic Service Change Acknowledge
21	2	DSD-REQ	Dynamic Service Deletion Request
22	2	DSD-RSP	Dynamic Service Deletion Response
23-255			Reserved for future use

RSVD 1 octet. This field is used to align the message payload on a 32 bit boundary. Set to 0 for this version.

Management Message Payload Variable length. As defined for each specific management message.

CRC Covers message including header fields (DA, SA,...). Polynomial defined by [ISO8802-3].

A compliant CMTS or CM **MUST** support the MAC management message types listed in Table 6-17, except messages specific to Telephony Return devices which **MAY** be supported.

6.3.2 Time Synchronization (SYNC)

Time Synchronization (SYNC) **MUST** be transmitted by CMTS at a periodic interval to establish MAC sublayer timing. This message **MUST** use an FC field with FC_TYPE = MAC Specific Header and FC_PARM = Timing MAC Header. This **MUST** be followed by a Packet PDU in the format shown in Figure 6-15.

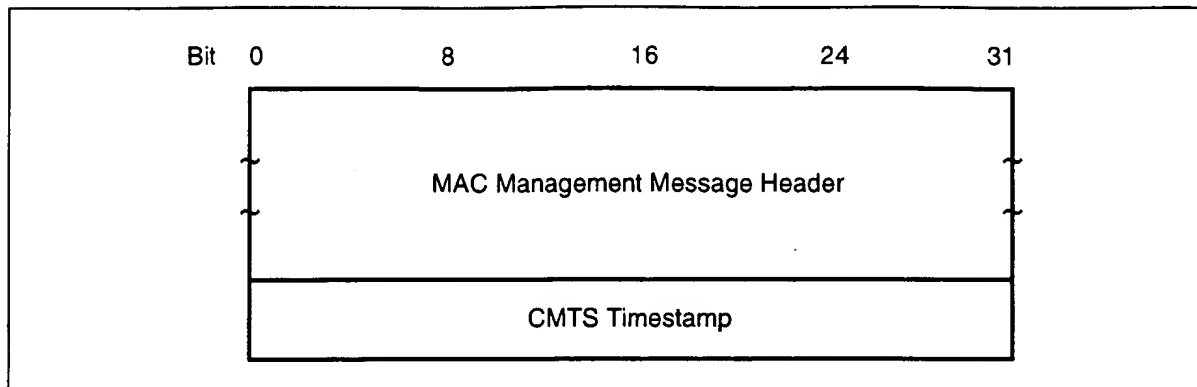


Figure 6-15. Format of Packet PDU Following the Timing Header

The parameters shall be as defined below.

CMTS Timestamp The count state of an incrementing 32 bit binary counter clocked with the CMTS 10.24 MHz master clock.

The CMTS timestamp represents the count state at the instant that the first byte (or a fixed time offset from the first byte) of the Time Synchronization MAC Management Message is transferred from the Downstream Transmission Convergence Sublayer to the Downstream Physical Media Dependent Sublayer as described in Section 4.3.7. The CMTS **MUST NOT** allow a SYNC message to cross an MPEG packet boundary¹.

1. Since the SYNC message applies to all upstream channels within this MAC domain, units were chosen to be independent of the symbol rate of any particular upstream channel. A timebase tick represents one half the smallest possible mini-slot at the highest possible symbol rate. See Section 7.3.4 for time-unit relationships.

6.3.3 Upstream Channel Descriptor (UCD)

An Upstream Channel Descriptor **MUST** be transmitted by the CMTS at a periodic interval to define the characteristics of an upstream channel (Figure 6-16). A separate message **MUST** be transmitted for each active upstream.

To provide for flexibility the message parameters following the channel ID **MUST** be encoded in a type/length/value (TLV) form in which the type and length fields are each 1 octet long.

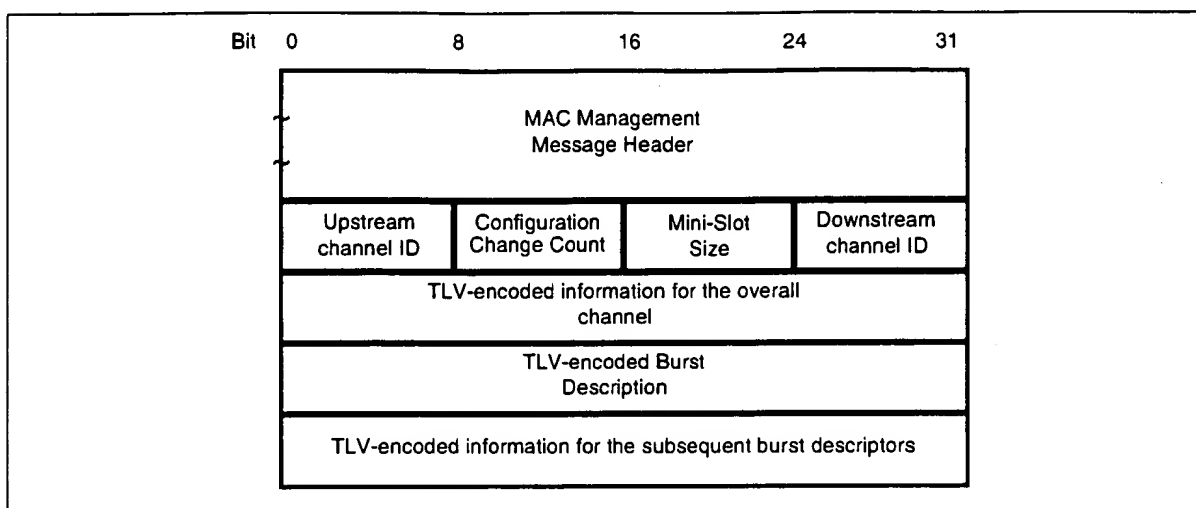


Figure 6-16. Upstream Channel Descriptor

A CMTS **MUST** generate UCDs in the format shown in Figure 6-16, including all of the following parameters:

- | | |
|-----------------------------------|--|
| Configuration Change Count | Incremented by one (modulo the field size) by the CMTS whenever any of the values of this channel descriptor change. If the value of this count in a subsequent UCD remains the same, the CM can quickly decide that the remaining fields have not changed, and may be able to disregard the remainder of the message. This value is also referenced from the MAP. |
| Mini-Slot Size | The size T of the Mini-Slot for this upstream channel in units of the Timebase Tick of $6.25 \mu\text{s}$. Allowable values are $T = 2^M$, $M = 1, \dots, 7$. That is, $T = 2, 4, 8, 16, 32, 64$ or 128 . |
| Upstream Channel ID | The identifier of the upstream channel to which this message refers. This identifier is arbitrarily chosen by the CMTS and is only unique within the MAC-Sublayer domain. |
| Downstream Channel ID | The identifier of the downstream channel on which this message has been transmitted. This identifier is arbitrarily chosen by the CMTS and is only unique within the MAC-Sublayer domain. |

All other parameters are coded as TLV tuples. The type values used **MUST** be those defined in Table 6-18, for channel parameters, and Table 6-19, for upstream physical layer burst attributes. Channel-wide parameters (types 1-3 in Table 6-18) **MUST** precede burst descriptors (type 4 below).

Table 6-18. Channel TLV Parameters

Name	Type (1 byte)	Length (1 byte)	Value (Variable length)
Symbol Rate	1	1	Multiples of base rate of 160 ksym/sec. (Value is 1, 2, 4, 8, or 16.)
Frequency	2	4	Upstream center frequency (Hz)
Preamble Pattern	3	1-128	Preamble superstring. All burst-specific preamble values are chosen as bit-substrings of this string. The first byte of the Value field contains the first 8 bits of the superstring, with the first bit of the preamble superstring in the MSB position of the first Value field byte, the eighth bit of the preamble superstring in the LSB position of the first Value field byte; the second byte in the Value field contains the second eight bits of the superstring, with the ninth bit of the superstring in the MSB of the second byte and sixteenth bit of the preamble superstring in the LSB of the second byte, and so forth.
Burst Descriptor	4		May appear more than once; described below. The length is the number of bytes in the overall object, including embedded TLV items.

Burst Descriptors are compound TLV encodings that define, for each type of upstream usage interval, the physical-layer characteristics that are to be used during that interval. The upstream interval usage codes are defined in the MAP message (see Section 6.3.4 and Table 6-20).




Figure 6-17. Top-Level Encoding for a Burst Descriptor

A Burst Descriptor **MUST** be included for each Interval Usage Code that is to be used in the allocation MAP. The Interval Usage Code **MUST** be one of the values from Table 6-20.

Within each Burst Descriptor is an unordered list of Physical-layer attributes, encoded as TLV values. These attributes are shown in Table 6-19.

Table 6-19. Upstream Physical-Layer Burst Attributes

Name	Type (1 byte)	Length (1 byte)	Value (Variable Length)
Modulation Type	1	1	1 = QPSK, 2 = 16QAM
Differential Encoding	2	1	1 = on, 2 = off
Preamble Length	3	2	Up to 1024 bits. The value must be an integral number of symbols (a multiple of 2 for QPSK and 4 for 16QAM)
Preamble Value Offset	4	2	Identifies the bits to be used for the preamble value. This is specified as a starting offset into the Preamble Pattern (see Table 6-18). That is, a value of zero means that the first bit of the preamble for this burst type is the value of the first bit of the Preamble Pattern. A value of 100 means that the preamble is to use the 101st and succeeding bits from the Preamble Pattern. This value must be a multiple of the symbol size. The first bit of the Preamble Pattern is the first bit into the symbol mapper (Figure 4-8), and is 11 in the first symbol of the burst (see Section 4.2.2.2).
FEC Error Correction (T)	5	1	0-10 (0 implies no FEC. The number of codeword parity bytes is 2^T)
FEC Codeword Information Bytes (k)	6	1	Fixed: 16 to 253 (assuming FEC on) Shortened: 16 to 253 (assuming FEC on) (Not used if no FEC, T=0)
Scrambler Seed	7	2	The 15-bit seed value left justified in the 2 byte field. Bit 15 is the MSB of the first byte and the LSB of the second byte is not used. (Not used if scrambler is off)
Maximum Burst Size	8	1	The maximum number of mini-slots that can be transmitted during this burst type. Absence of this configuration setting implies that the burst size is limited elsewhere. When the interval type is Short Data Grant this value MUST be present and greater than zero. (See 7.1.2.5)
Guard Time Size	9	1	Number of symbol times which must follow the end of this burst. (Although this value may be derivable from other network and architectural parameters, it is included here to ensure that the CMs and CMTS all use the same value.)
Last Codeword Length	10	1	1 = fixed; 2 = shortened
Scrambler on/off	11	1	1 = on; 2 = off

6.3.3.1 Example of UCD Encoded TLV Data

An example of UCD encoded TLV data is given in Figure 6-18.

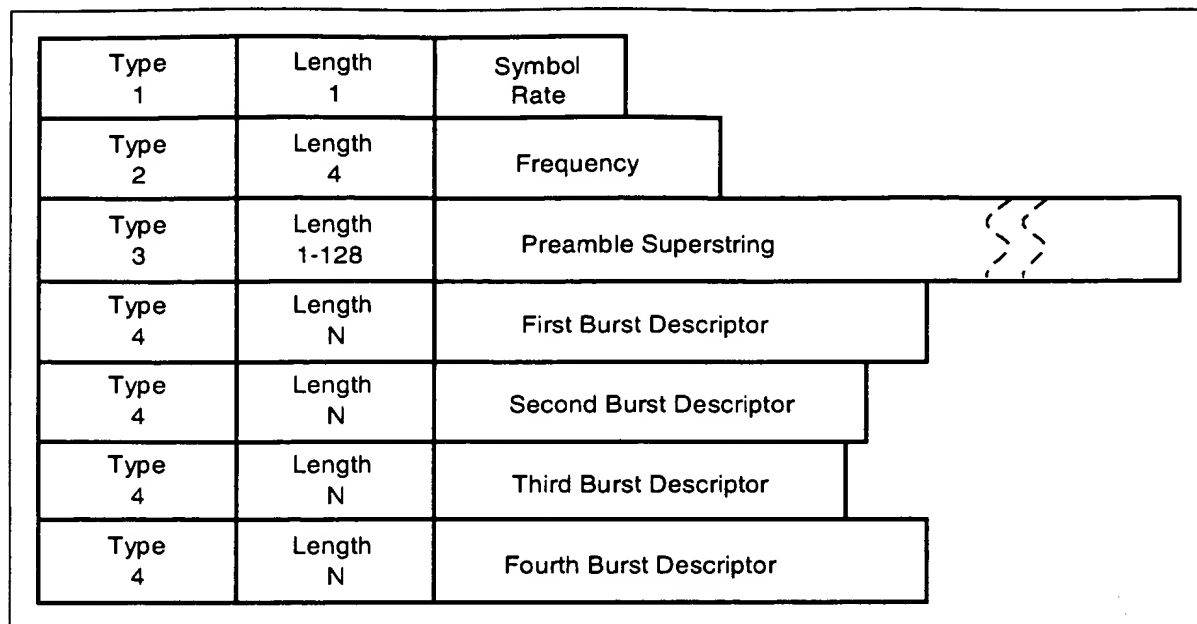


Figure 6-18. Example of UCD Encoded TLV Data

6.3.4 Upstream Bandwidth Allocation Map (MAP)

A CMTS MUST generate MAPs in the format shown in Figure 6-19.




Figure 6-19. MAP Format

The parameters MUST be as follows:

Upstream Channel ID	The identifier of the upstream channel to which this message refers.
UCD Count	Matches the value of the Configuration Change Count of the UCD which describes the burst parameters which apply to this map. See Section 9.3.2.
Number Elements	Number of information elements in the map.
Reserved	Reserved field for alignment.
Alloc Start Time	Effective start time from CMTS initialization (in mini-slots) for assignments within this map.
Ack Time	Latest time, from CMTS initialization, (mini-slots) processed in upstream. This time is used by the CMs for collision detection purposes. See Section 7.4.
Ranging Backoff Start	Initial back-off window for initial ranging contention, expressed as a power of two. Values range 0-15 (the highest order bits must be unused and set to 0).
Ranging Backoff End	Final back-off window for initial ranging contention, expressed as a power of two. Values range 0-15 (the highest order bits must be unused and set to 0).

Data Backoff Start Initial back-off window for contention data and requests, expressed as a power of two. Values range 0-15 (the highest order bits must be unused and set to 0).

Data Backoff End Final back-off window for contention data and requests, expressed as a power of two. Values range 0-15 (the highest order bits must be unused and set to 0).

MAP Information Elements MUST be in the format defined in Figure 6-20 and Table 6-20. Values for IUCs are defined in Table 6-20 and are described in detail in Section 7.1.2.

Note: That the lower (26-M) bits of the Alloc Start Time and Ack Time MUST be used as the effective MAP start and ack times where M is given in Section 6.3.3. The relationship between the Alloc Start/Ack time counters and the timestamp counter is described in Section 7.4.

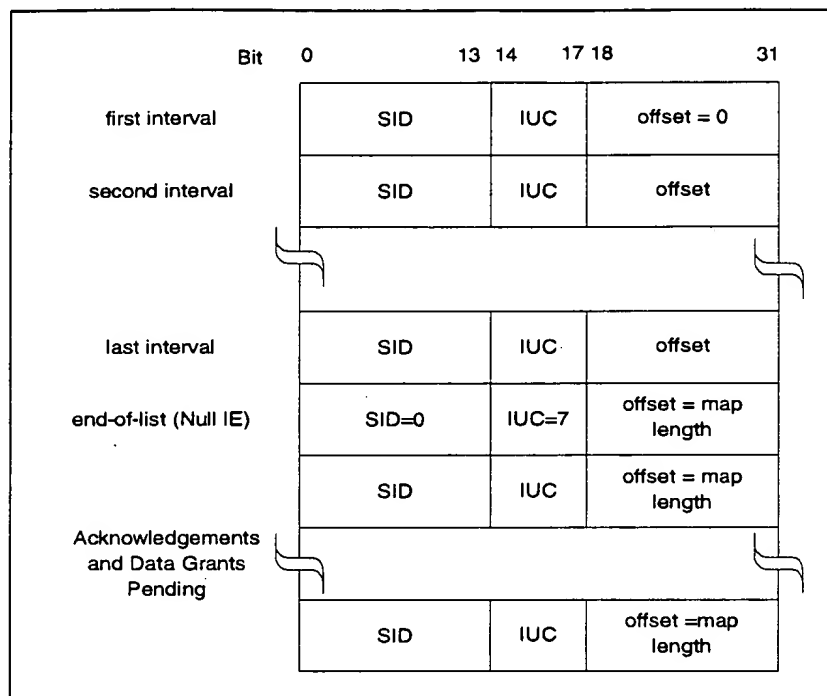


Figure 6-20. MAP Information Element Structure

Table 6-20. Allocation MAP Information Elements (IE)

IE Name ^a	Interval Usage Code (IUC) (4 bits)	SID (14 bits)	Mini-slot Offset (14 bits)
Request	1	any	Starting offset of REQ region
REQ/Data (refer to Appendix A for multicast definition)	2	multicast	Starting offset of IMMEDIATE Data region (well-known multicasts define start intervals)
Initial Maintenance	3	broadcast/ multicast	Starting offset of MAINT region (used in Initial Ranging)
Station Maintenance ^b	4	unicast ^c	Starting offset of MAINT region (used in Periodic Ranging)
Short Data Grant ^d	5	unicast	Starting offset of Data Grant assignment; If inferred length = 0, then it is a Data Grant pending.
Long Data Grant	6	unicast	Starting offset of Data Grant assignment; If inferred length = 0, then it is a Data Grant Pending
Null IE	7	zero	Ending offset of the previous grant. Used to bound the length of the last actual interval allocation.
Data Ack	8	unicast	CMTS sets to map length
Reserved	9-14	any	Reserved
Expansion	15	expanded IUC	# of additional 32-bit words in this IE

- Each IE is a 32-bit quantity, of which the most significant 14 bits represent the SID, the middle 4 bits the IUC, and the low-order 14 bits the mini-slot offset.
- Although the distinction between Initial Maintenance and Station Maintenance is unambiguous from the Service ID type, separate codes are used to ease physical-layer configuration (see burst descriptor encodings, Table 6-19).
- The SID used in the Station Maintenance IE MUST be a Temporary SID, or the first Registration SID (and maybe the only one) that was assigned in the REG-RSP message to a CM.
- The distinction between long and short data grants is related to the amount of data that can be transmitted in the grant. A short data grant interval may use FEC parameters that are appropriate to short packets while a long data grant may be able to take advantage of greater FEC coding efficiency.

6.3.5 Ranging Request (RNG-REQ)

A Ranging Request **MUST** be transmitted by a CM at initialization and periodically on request from CMTS to determine network delay and request power adjustment. This message **MUST** use an FC_TYPE = MAC Specific Header and FC_PARM = Timing MAC Header. This **MUST** be followed by a Packet PDU in the format shown in Figure 6-21.




Figure 6-21. Packet PDU Following the Timing Header

Parameters **MUST** be as follows:

SID

For RNG-REQ messages transmitted in Initial Maintenance intervals:

- Initialization SID if modem is attempting to join the network
- Initialization SID if modem has not yet registered and is changing downstream (or both downstream and upstream) channels as directed by a downloaded parameter file
- Temporary SID if modem has not yet registered and is changing upstream (not downstream) channels as directed by a downloaded parameter file
- Registration SID (previously assigned in REG-RSP) if modem is registered and is changing upstream channels

For RNG-REQ messages transmitted in Station Maintenance intervals:

- Assigned SID

This is a 16-bit field of which the lower 14 bits define the SID with bits 14,15 defined to be 0.

Downstream Channel ID

The identifier of the downstream channel on which the CM received the UCD which described this upstream. This is an 8-bit field.

Pending Till Complete

If zero, then all previous Ranging Response attributes have been applied prior to transmitting this request. If nonzero then this is time estimated to be needed to complete assimilation of ranging parameters. Note that only equalization can be deferred. Units are in unsigned centiseconds (10 msec).

6.3.6 Ranging Response (RNG-RSP)

A Ranging Response **MUST** be transmitted by a CMTS in response to received RNG-REQ. The state machines describing the ranging procedure appear in Section 9.2.4. In that procedure it may be noted that, from the point of view of the CM, reception of a Ranging Response is stateless. In particular, the CM **MUST** be prepared to receive a Ranging Response at any time, not just following a Ranging Request.

To provide for flexibility, the message parameters following the Upstream Channel ID **MUST** be encoded in a type/length/value (TLV) form.




Figure 6-22. Ranging Response

A CMTS **MUST** generate Ranging Responses in the form shown in 6-22, including all of the following parameters:

SID	If the modem is being instructed by this response to move to a different channel, this is initialization SID. Otherwise, this is the SID from the corresponding RNG-REQ to which this response refers, except that if the corresponding RNG-REQ was an initial ranging request specifying a initialization SID, then this is the assigned temporary SID.
Upstream Channel ID	The identifier of the upstream channel on which the CMTS received the RNG-REQ to which this response refers.
Ranging Status	Used to indicate whether upstream messages are received within acceptable limits by CMTS.
All other parameters are coded as TLV tuples.	
Timing Adjust Information	The time by which to offset frame transmission so that frames arrive at the expected mini-slot time at the CMTS.
Power Adjust Information	Specifies the relative change in transmission power level that the CM is to make in order that transmissions arrive at the CMTS at the desired power.
Frequency Adjust Information	Specifies the relative change in transmission frequency that the CM is to make in order to better match the CMTS. (This is fine-frequency adjustment within a channel, not re-assignment to a different channel)

CM Transmitter Equalization Information If the CM implements transmission equalization, this provides the equalization coefficients (optional).

Downstream Frequency Override An optional parameter. The downstream frequency with which the modem should redo initial ranging. (See Section 6.3.6.3)

Upstream Channel ID Override An optional parameter. The identifier of the upstream channel with which the modem should redo initial ranging. (See Section 6.3.6.3)

6.3.6.1 Encodings

The type values used **MUST** be those defined in Table 6-21 and Figure 6-23. These are unique within the ranging response message but not across the entire MAC message set. The type and length fields **MUST** each be 1 octet in length.

Table 6-21. Ranging Response Message Encodings

Name	Type (1 byte)	Length (1 byte)	Value (Variable Length)
Timing Adjust	1	4	TX timing offset adjustment (signed 32-bit, units of (6.25 microsec/64))
Power Level Adjust	2	1	TX Power offset adjustment (signed 8-bit, 1/4-dB units)
Offset Frequency Adjust	3	2	TX frequency offset adjustment (signed 16-bit, Hz units)
Transmit Equalization Adjust	4	n	TX equalization data - see details below
Ranging Status	5	1	1 = continue, 2 = abort, 3 = success
Downstream frequency override	6	4	Center frequency of new downstream channel in Hz
Upstream channel ID override	7	1	Identifier of the new upstream channel.
Reserved	8-255	n	Reserved for future use

type 4	length	Number of taps per symbol
number of forward taps (N)	number of reverse taps (M)	
first coefficient F_0 (real)		first coefficient F_0 (imaginary)
last coefficient F_N (real)		last coefficient F_N (imaginary)
first reverse coefficient D_0 (real)		first reverse coefficient D_0 (imaginary)
last reverse coefficient D_M (real)		last reverse coefficient D_M (imaginary)

Figure 6-23. Generalized Decision Feedback Equalization Coefficients

The total number of taps per symbol **MUST** be in the range 1 to 4. The total number of taps **MAY** range up to 64. Each tap consists of a real and imaginary coefficient entry in the table.

If more than 255 bytes are needed to represent equalization information, then several type-4 elements **MAY** be used. Data **MUST** be treated as if byte-concatenated, that is, the first byte after the length field of the second type-4 element is treated as if it immediately followed the last byte of the first type-4 element.

The coefficients that are sent to the CM may be coefficients of a CMTS demodulator equalizer such as shown in Figure 6-24, which, after acquisition, will have tap values that represent the channel distortion. Other equalization methods may be devised in the future. If so, they will use a different type-value so that the element is not overloaded. This is a vendor-specific issue which is not described here.

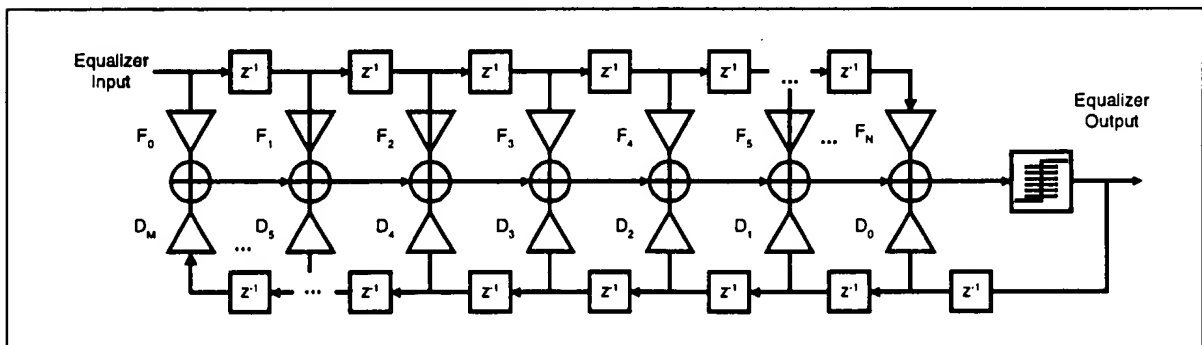


Figure 6-24. CMTS Demodulator Equalizer Tap Location Definition

6.3.6.2 Example of TLV Data

An example of TLV data is given in Figure 6-25.

Type 1	Length 4	Timing adjust	
Type 2	Length 1	Power adjust	
Type 3	Length 2	Frequency adjust information	
Type 4	Length x	x bytes of CM transmitter equalization information	
Type 5	Length 1	Ranging status	

Figure 6-25. Example of TLV Data

6.3.6.3 Overriding Channels During Initial Ranging

The RNG-RSP message allows the CMTS to instruct the modem to move to a new downstream and/or upstream channel and to repeat initial ranging. However, the CMTS may do this only in response to an initial ranging request from a modem that is attempting to join the network, or in response to any of the unicast ranging requests that take place immediately after this initial ranging and up to the point where the modem successfully completes periodic ranging. If a downstream frequency override is specified in the RNG-RSP, the modem **MUST** reinitialize its MAC and perform initial ranging using the specified downstream center frequency as the first scanned channel. For the upstream channel, the modem may select any valid channel based on received UCD messages.

If an upstream channel ID override is specified in the RNG-RSP, the modem **MUST** reinitialize its MAC and perform initial ranging using for its first attempt the upstream channel specified in the RNG-RSP and the same downstream frequency on which the RNG-RSP was received.

If both downstream frequency and upstream channel ID overrides are present in the RNG-RSP, the modem **MUST** reinitialize its MAC and perform initial ranging using for its first attempt the specified downstream frequency and upstream channel ID.

Note that when a modem with an assigned temporary SID is instructed to move to a new downstream and/or upstream channel and to redo initial ranging, the modem **MUST** consider the temporary SID to be deassigned. The modem **MUST** redo initial ranging using the Initialization SID.

Configuration file settings for upstream channel ID and downstream frequency are optional, but if specified in the config file they take precedence over the ranging response parameters. Once ranging is complete, only the C.1.1.2 and UCC-REQ mechanisms are available for moving the modem to a new upstream channel, and only the C.1.1.1 mechanism is available for moving the modem to a new downstream channel.

6.3.7 Registration Request (REG-REQ)

A Registration Request **MUST** be transmitted by a CM at initialization after receipt of a CM parameter file.

To provide for flexibility, the message parameters following the SID **MUST** be encoded in a type/length/value form.




Figure 6-26. Registration Request

A CM **MUST** generate Registration Requests in the form shown in Figure 6-26, including the following parameters:

SID Temporary SID for this CM.¹

All other parameters are coded as TLV tuples as defined in Appendix C.

Registration Requests can contain many different TLV parameters, some of which are set by the CM according to its configuration file and some of which are generated by the CM itself. If found in the Configuration File, the following Configuration Settings **MUST** be included in the Registration Request.

Configuration File Settings:

- Downstream Frequency Configuration Setting
- Upstream Channel ID Configuration Setting
- Network Access Control Object
- Upstream Packet Classification Configuration Setting
- Downstream Packet Classification Configuration Setting
- Class of Service Configuration Setting
- Upstream Service Flow Configuration Setting
- Downstream Service Flow Configuration Setting
- Baseline Privacy Configuration Setting
- Maximum Number of CPEs

1. edited per rfi-n-99043 06/21/99

- Maximum Number of Classifiers
- Privacy Enable Configuration Setting
- Payload Header Suppression
- TFTP Server Timestamp
- TFTP Server Provisioned Modem Address
- Vendor-Specific Information Configuration Setting
- CM MIC Configuration Setting
- CMTS MIC Configuration Setting

Note: The CM MUST forward the vendor specific configuration settings to the CMTS in the same order in which they were received in the configuration file to allow the message integrity check to be performed.

The following registration parameter MUST be included in the Registration Request.

Vendor Specific Parameter:

- Vendor ID Configuration Setting (Vendor ID of CM)

The following registration parameter MUST also be included in the Registration Request.

- Modem Capabilities Encodings¹

The following registration parameter MAY also be included in the Registration Request.

- Modem IP Address

The following Configuration Settings MUST NOT be forwarded to the CMTS in the Registration Request.

- Software Upgrade Filename
- Software Upgrade TFTP Server IP Address
- SNMP Write-Access Control
- SNMP MIB Object
- CPE Ethernet MAC Address
- HMAC Digest
- End Configuration Setting
- Pad Configuration Setting
- Telephone Settings Option

1. The CM MUST specify all of its Modem Capabilities in its Registration Request. The CMTS MUST NOT assume any Modem Capability which is defined but not explicitly indicated in the CM's Registration Request.

6.3.8 Registration Response (REG-RSP)

A Registration Response **MUST** be transmitted by CMTS in response to received REG-REQ.

To provide for flexibility, the message parameters following the Response field **MUST** be encoded in a TLV format.




Figure 6-27. Registration Response Format

A CMTS **MUST** generate Registration Responses in the form shown in Figure 6-27, including both of the following parameters:

SID from Corresponding REG-REQ

SID from corresponding REG-REQ to which this response refers. (This acts as a transaction identifier)

Response

0 = Okay
1 = Authentication Failure
2 = Class of Service Failure

Note: Failures apply to the entire Registration Request. Even if only a single requested Service Flow or DOCSIS 1.0 Service Class is invalid or undeliverable the entire registration is failed.

If the REG-REQ was successful, and contained Service Flow Parameters, Classifier Parameters, or Payload Header Suppression Parameters, the REG-RSP **MUST** contain, for each of these:

Classifier Parameters

All of the Classifier Parameters from the corresponding REG-REQ, plus the Classifier Identifier assigned by the CMTS.

Service Flow Parameters All the Service Flow Parameters from the REG-REQ, plus the Service Flow ID assigned by the CMTS. Every Service Flow that contained a Service Class Name that was admitted/activated¹ MUST be expanded into the full set of TLVs defining the Service Flow. Every upstream Service Flow that was admitted/activated MUST have a Service Identifier assigned by the CMTS. A Service Flow that was only provisioned will include only those QoS parameters that appeared in the REG-REQ, plus the assigned Service Flow ID.

Payload Header Suppression Parameters

All the Payload Header Suppression Parameters from the REG-REQ, plus the Payload Header Suppression Index assigned by the CMTS.

If the REG-REQ failed, and contained Service Flow Parameters, Classifier Parameters, or Payload Header Suppression Parameters, the REG-RSP MUST contain at least one of the following:

Classifier Error Set A Classifier Error Set and identifying Classifier Reference and Service Flow Reference MUST be included for every failed Classifier in the corresponding REG-REQ. Every Classifier Error Set MUST include every specific failed Classifier Parameter of the corresponding Classifier.

Service Flow Error Set A Service Flow Error Set and identifying Service Flow Reference MUST be included for every failed Service Flow in the corresponding REG-REQ. Every Service Flow Error Set MUST include every specific failed QoS Parameter of the corresponding Service Flow.

Payload Header Suppression Error Set

A PHS Error Set and identifying Classifier Reference MUST be included for every failed PHS Rule in the corresponding REG-REQ. Every PHS Error Set MUST include every specific failed PHS Parameter of the corresponding failed PHS Rule.

Service Class Name expansion always occurs at admission time. Thus, if a Registration-Request contains a Service Flow Reference and a Service Class Name for deferred admission/activation, the Registration-Response MUST NOT include any additional QoS Parameters except the Service Flow Identifier. (Refer to Section 8.1.3)

If the corresponding Registration Request contains DOCSIS 1.0 Service Class TLV's (refer to C.1.1.4), the Registration Response MUST contain the following TLV tuples:

DOCSIS 1.0 Service Class Data Returned when Response = Okay

Service ID / service class tuple for each class of service granted

Note: Service class IDs MUST be those requested in the corresponding REG-REQ.

Service Not Available

Returned when Response = Class of Service Failure.

If a service class cannot be supported, this configuration setting is returned in place of the service class data.

All other parameters are coded TLV tuples

Modem Capabilities

The CMTS response to the capabilities of the modem (if present in the Registration Request)

1. The ActiveQoSParamSet or AdmittedQoSParamSet is non-null.

Vendor-Specific Data

As defined in Appendix C

- Vendor ID Configuration Setting (vendor ID of CMTS)
- Vendor-specific extensions

Note: The temporary SID MUST no longer be used once the REG-RSP is received.

6.3.8.1 Encodings

The type values used MUST be those shown below. These are unique within the Registration Response message but not across the entire MAC message set. The type and length fields MUST each be 1 octet.

6.3.8.1.1 Modem Capabilities

This field defines the CMTS response to the modem capability field in the Registration Request. The CMTS responds to the modem capabilities to indicate whether they may be used. If the CMTS does not recognize a modem capability, it must return this as "off" in the Registration Response.

Only capabilities set to "on" in the REG-REQ may be set "on" in the REG-RSP as this is the handshake indicating that they have been successfully negotiated.

Encodings are as defined for the Registration Request.

6.3.8.1.2 DOCSIS 1.0 Service Class Data

A DOCSIS 1.0 Service Class Data parameter MUST be present in the Registration Response for each DOCSIS 1.0 Class of Service parameter (refer to C.1.1.4) in the Registration Request.

This encoding defines the parameters associated with a requested class of service. It is somewhat complex in that it is composed from a number of encapsulated type/length/value fields. The encapsulated fields define the particular class of service parameters for the class of service in question. Note that the type fields defined are only valid within the encapsulated service class data configuration setting string. A single service class data configuration setting MUST be used to define the parameters for a single service class. Multiple class definitions MUST use multiple service class data configuration setting sets.

Each received DOCSIS 1.0 Class of Service parameter must have a unique Class ID in the range 1..16. If no Class ID was present for any single DOCSIS 1.0 Class-of-Service TLV in the REG-REQ, the CMTS MUST send a REG-RSP with a class-of-service failure response and no DOCSIS 1.0 Class-of-Service TLVs.

Type	Length	Value
1	n	Encoded service class data

Class ID

The value of the field MUST specify the identifier for the class of service to which the encapsulated string applies. This MUST be a class which was requested in the associated REG-REQ, if present.

Type	Length	Value
1	1	from REG-REQ

Valid Range

The class ID MUST be in the range 1 to 16.

Service ID

The value of the field **MUST** specify the SID associated with this service class.

Type	Length	Value
2	2	SID

6.3.9 Registration Acknowledge (REG-ACK)

A Registration Acknowledge **MUST** be transmitted by the CM in response to a REG-RSP from the CMTS.¹ It confirms acceptance by the CM of the QoS parameters of the flow as reported by the CMTS in its REG-RSP. The format of a REG-ACK **MUST** be as shown in Figure 6-28.

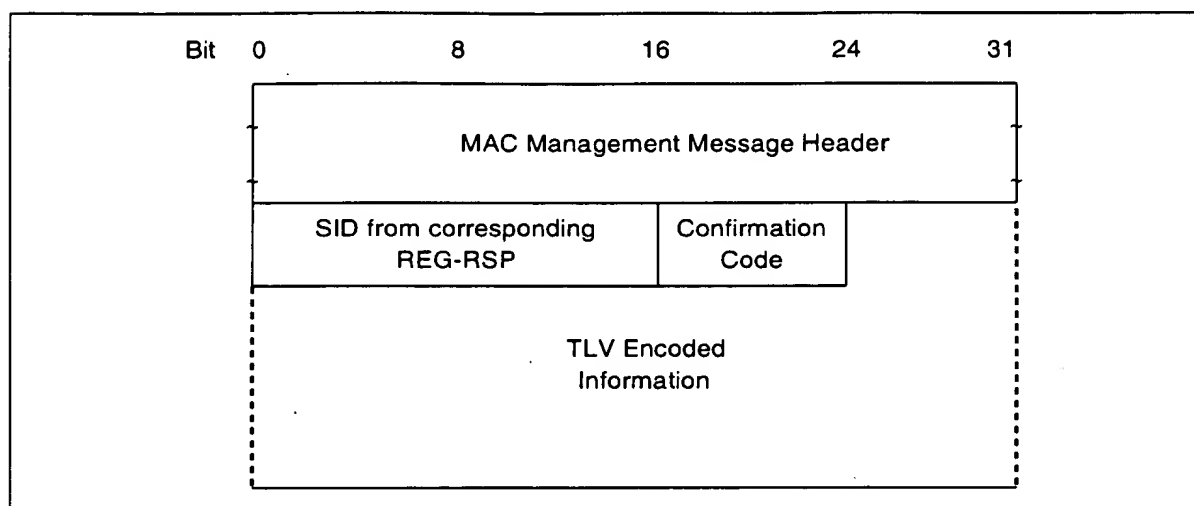


Figure 6-28. Registration Acknowledgment

The parameter **MUST** be as follows:

SID from Corresponding REG-RSP

SID from corresponding REG-RSP to which this acknowledgment refers. (This acts as a transaction identifier)

Confirmation Code

The appropriate Confirmation Code (refer to C.4) for the entire corresponding Registration Response.

The CM **MUST** forward all provisioned Classifiers, Service Flows and Payload Header Suppression Rules to the CMTS. Since any of these provisioned items can fail, the REG-ACK **MUST** include Error Sets for all failures related to these provisioned items.

Classifier Error Set

A Classifier Error Set and identifying Classifier Identifier and Service Flow Identifier pair **MUST** be included for every failed Classifier in the corresponding REG-RSP. Every Classifier Error Set **MUST** include every specific failed Classifier Parameter of the corresponding failed Classifier. This parameter **MUST** be omitted if the entire REG-REQ/RSP is successful.

1. The Registration-Acknowledge is a DOCSIS 1.1 message. Refer to Appendix G for details of registration interoperability issues.

Service Flow Error Set

The Service Flow Error Set of the REG-ACK message encodes specifics of any failed Service Flows in the REG-RSP message. A Service Flow Error Set and identifying Service Flow Reference **MUST** be included for every failed QoS Parameter of every failed Service Flow in the corresponding REG-RSP message. This parameter **MUST** be omitted if the entire REG-REQ/RSP is successful.

Payload Header Suppression Error Set

A PHS Error Set and identifying PHS Index and Classifier Reference/Identifier pair **MUST** be included for every failed PHS Rule in the corresponding REG-RSP. Every PHS Error Set **MUST** include every specific failed PHS Parameter of the corresponding failed PHS Rule. This parameter **MUST** be omitted if the entire REG-REQ/RSP is successful.

Note: Per Service Flow acknowledgment is necessary not just for synchronization between the CM and CMTS, but also to support use of the Service Class Name. (Refer to Section 8.1.3) Since the CM may not know all of the Service Flow parameters associated with a Service Class Name when making the Registration Request, it may be necessary for the CM to NAK a Registration Response if it has insufficient resources to actually support this Service Flow.

6.3.10 Upstream Channel Change Request (UCC-REQ)

An Upstream Channel Change Request MAY be transmitted by a CMTS to cause a CM to change the upstream channel on which it is transmitting. The format of an UCC-REQ message is shown in Figure 6-29.




Figure 6-29. Upstream Channel Change Request

Parameters MUST be as follows:

Upstream Channel ID The identifier of the upstream channel to which the CM is to switch for upstream transmissions. This is an 8-bit field.

All other parameters are coded as TLV tuples.

Ranging Technique Directions for the type of ranging that the CM should perform once synchronized to the new upstream channel.

6.3.10.1 Encodings

The type values used MUST be those shown below. These are unique within the Upstream Channel Change Request message, but not across the entire MAC message set. The type and length fields MUST each be 1 octet.

6.3.10.1.1 Ranging Technique

When present, this TLV allows the CMTS to direct the CM what level of re-ranging, if any, to perform. The CMTS can make this decision based upon its knowledge of the differences between the old and new upstream channels.

For example, areas of upstream spectrum are often configured in groups. A UCC-REQ to an adjacent channel within a group may not warrant re-ranging. Alternatively, a UCC-REQ to a non-adjacent channel might require station maintenance whereas a UCC-REQ from one channel group to another might require initial maintenance.

Type	Length	Value
1	1	0 = Perform initial maintenance on new channel 1 = Perform only station maintenance on new channel 2 = Perform either initial maintenance or station maintenance on new channel ¹ 3 = Use the new channel directly without performing initial or station maintenance

If this TLV is absent, the CM MUST perform ranging with initial maintenance. For backwards compatibility, the CMTS MUST accept a CM which ignores this tuple and performs initial maintenance.

Note: This option should not be used in physical plants where upstream transmission characteristics are not consistent.

6.3.11 Upstream Channel Change Response (UCC-RSP)

An Upstream Channel Change Response MUST be transmitted by a CM in response to a received Upstream Channel Change Request message to indicate that it has received and is complying with the UCC-REQ. The format of an UCC-RSP message is shown in Figure 6-30.

Before it begins to switch to a new upstream channel, a CM MUST transmit a UCC-RSP on its existing upstream channel. A CM MAY ignore an UCC-REQ message while it is in the process of performing a channel change. When a CM receives a UCC-REQ message requesting that it switch to an upstream channel that it is already using, the CM MUST respond with a UCC-RSP message on that channel indicating that it is already using the correct channel.

After switching to a new upstream channel, a CM MUST re-range using the Ranging Technique in the corresponding UCC-REQ, and then will proceed without re-performing registration. The full procedure for changing channels is described in Section 9.3.3.

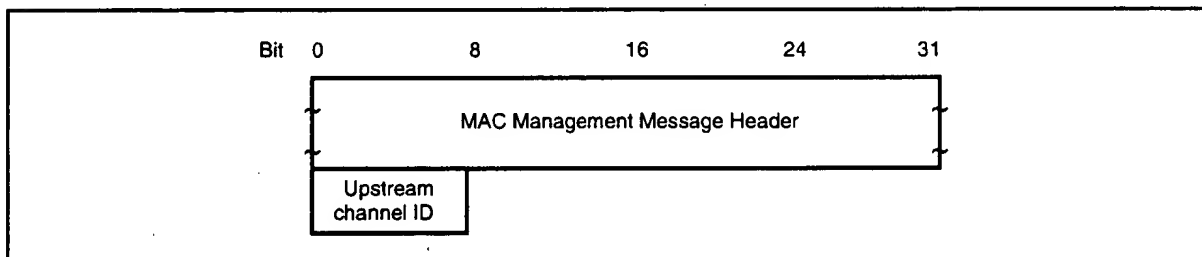


Figure 6-30. Upstream Channel Change Response

Parameters MUST be as follows:

Upstream Channel ID The identifier of the upstream channel to which the CM is to switch for upstream transmissions. This is the same Channel ID specified in the UCC-REQ message. This is an 8-bit field.

1. This value authorizes a CM to use an initial maintenance or station maintenance region, which ever occurs first. This value might be used when there is uncertainty when the CM may execute the UCC and thus a chance that it might miss station maintenance slots.

6.3.12 Dynamic Service Addition — Request (DSA-REQ)

A Dynamic Service Addition Request MAY be sent by a CM or CMTS to create a new Service Flow.

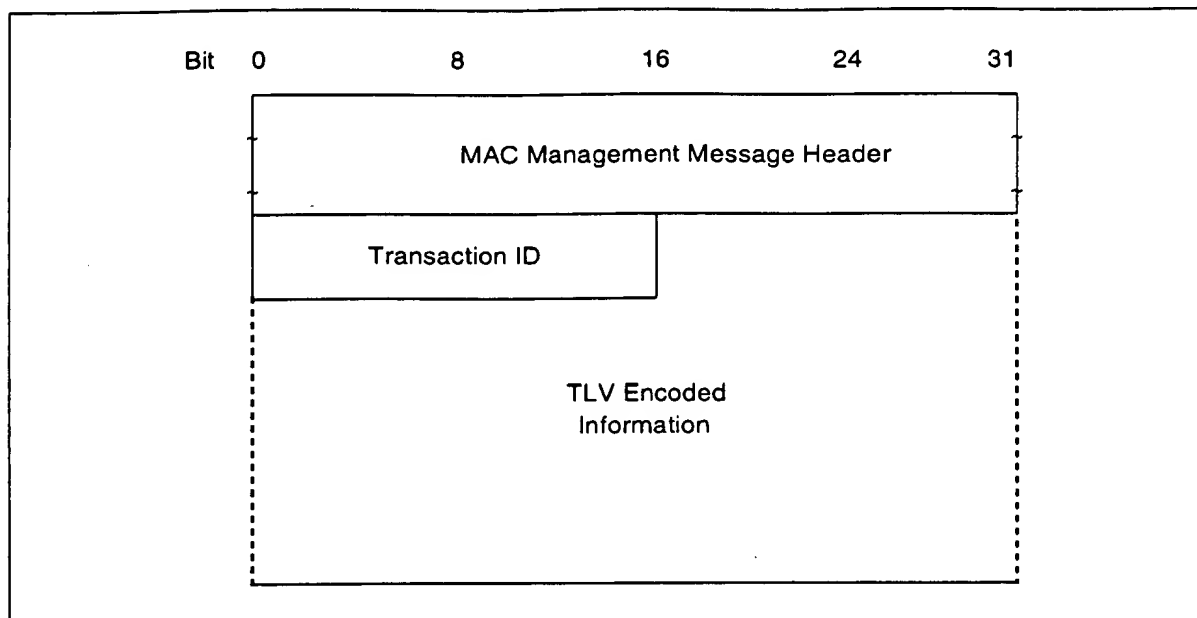


Figure 6-31. Dynamic Service Addition — Request

A CM or CMTS MUST generate DSA-REQ messages in the form shown in Figure 6-31 including the following parameter:

Transaction ID Unique identifier for this transaction assigned by the sender.

All other parameters are coded as TLV tuples as defined in Appendix C. A DSA-REQ message MUST NOT contain parameters for more than one Service Flow in each direction, i.e., a DSA-REQ message MUST contain parameters for either a single upstream Service Flow, or for a single downstream Service Flow, or for one upstream and one downstream Service Flow.¹

The DSA-REQ message MUST contain:

Service Flow Parameters Specification of the Service Flow's traffic characteristics and scheduling requirements.

The DSA-REQ message MAY contain classifier parameters and payload header suppression parameters associated with the Service Flows specified in the message:²

Classifier Parameters Specification of the rules to be used to classify packets into a specific Service Flow.

Payload Header Suppression

1. Paragraph edited per rfi-n-99048 06/30/99. ew

2. Paragraph edited per rfi-n-99048 06/30/99. ew

Parameters Specification of the payload header suppression rules to be used with an associated classifier.

If Privacy is enabled, the DSA-REQ message MUST contain:

HMAC-Digest The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute MUST be the final Attribute in the Dynamic Service message's Attribute list. (Refer to Appendix C.1.4.1)

6.3.12.1 CM-Initiated Dynamic Service Addition

CM-initiated DSA-Requests MUST use the Service Flow Reference to link Classifiers to Service Flows. Values of the Service Flow Reference are local to the DSA message; each Service Flow within the DSA-Request MUST be assigned a unique Service Flow Reference. This value need not be unique with respect to the other service flows known by the sender.

CM-initiated DSA-Request MUST use the Classifier Reference and Service Flow Reference to link Payload Header Suppression Parameters to Classifiers and Service Flows. Values of the Classifier Reference are many to one with Service Flows; each Classifier associated with a given Service Flow MUST be assigned a unique Classifier Reference.

CM-initiated DSA-Requests MAY use the Service Class Name (refer to C.2.2.3.4) in place of some, or all, of the QoS Parameters.

6.3.12.2 CMTS-Initiated Dynamic Service Addition

CMTS-initiated DSA-Requests MUST use the Service Flow ID to link Classifiers to Service Flows. Service Flow Identifiers are unique within the MAC domain. CMTS-initiated DSA-Requests for Upstream Service Flows MUST also include a Service ID.

CMTS-initiated DSA-Requests which include Classifiers, MUST assign a unique Classifier Identifier on a per Service Flow basis.

CMTS-initiated DSA-Requests for named Service Classes MUST include the QoS Parameter Set associated with that Service Class.

6.3.13 Dynamic Service Addition — Response (DSA-RSP)

A Dynamic Service Addition Response **MUST** be generated in response to a received DSA-Request. The format of a DSA-RSP **MUST** be as shown in Figure 6-32.

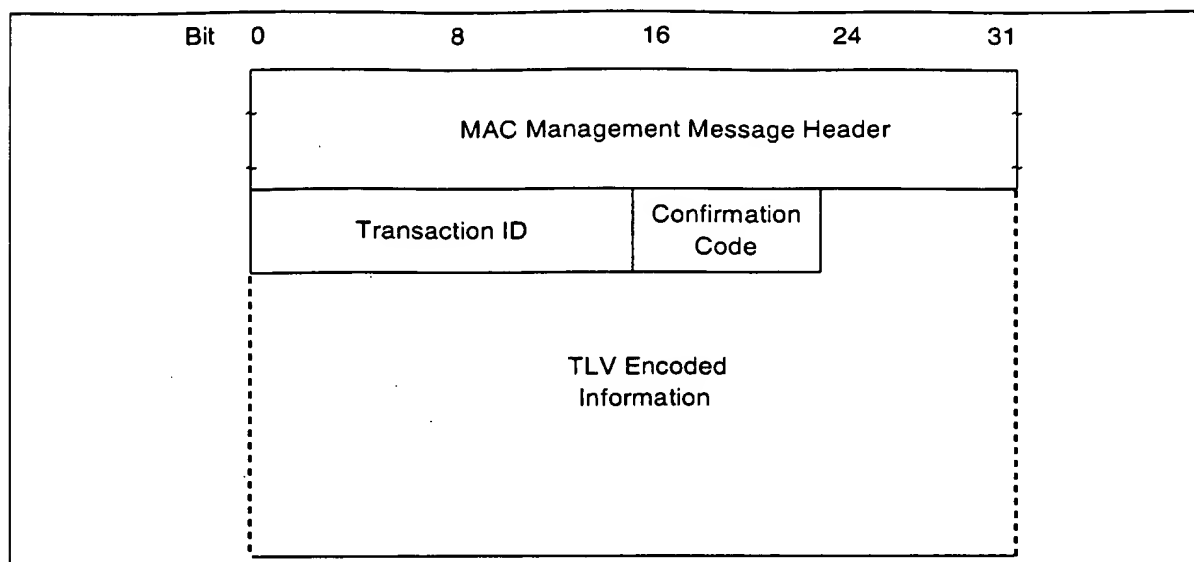


Figure 6-32. Dynamic Service Addition — Response

Parameters **MUST** be as follows:

Transaction ID	Transaction ID from corresponding DSA-REQ.
Confirmation Code	The appropriate Confirmation Code (refer to C.4) for the entire corresponding DSA-Request.

All other parameters are coded as TLV tuples as defined in Appendix C.

If the transaction is successful, the DSA-RSP **MAY** contain one or more of the following:

Classifier Parameters	The complete specification of the Classifier MUST be included in the DSA-RSP only if it includes a newly assigned Classifier Identifier. If a requested Classifier contained a Classifier Reference, the DSA-RSP MUST contain a Classifier Identifier.
Service Flow Parameters	The complete specification of the Service Flow MUST be included in the DSA-RSP only if it includes a newly assigned Service Flow Identifier or an expanded Service Class Name.
Payload Header Suppression Parameters	The complete specification of the PHS Parameters MUST be included in the DSA-RSP only if it includes a newly assigned PHS Index. If included, the PHS Parameters MUST contain a Classifier Identifier and a Service Flow Identifier.

If the transaction is unsuccessful, the DSA-RSP MUST include at least one of:

- Service Flow Error Set** A Service Flow Error Set and identifying Service Flow Reference/Identifier MUST be included for every failed Service Flow in the corresponding DSA-REQ message. Every Service Flow Error Set MUST include every specific failed QoS Parameter of the corresponding Service Flow. This parameter MUST be omitted if the entire DSA-REQ is successful.
- Classifier Error Set** A Classifier Error Set and identifying Classifier Reference/Identifier and Service Flow Reference/Identifier pair MUST be included for every failed Classifier in the corresponding DSA-REQ. Every Classifier Error Set MUST include every specific failed Classifier Parameter of the corresponding failed Classifier. This parameter MUST be omitted if the entire DSA-REQ is successful.
- Payload Header Suppression Error Set** A PHS Error Set and identifying Classifier Reference/Identifier and Service Flow Reference/Identifier pair MUST be included for every failed PHS Rule in the corresponding DSA-REQ. Every PHS Error Set MUST include every specific failed PHS Parameter of the corresponding failed PHS Rule. This parameter MUST be omitted if the entire DSA-REQ is successful.

If Privacy is enabled, the DSA-RSP message MUST contain:

- HMAC-Digest** The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute MUST be the final Attribute in the Dynamic Service message's Attribute list. (Refer to Appendix C.1.4.1)

6.3.13.1 CM-Initiated Dynamic Service Addition

The CMTS's DSA-Response for Service Flows that are successfully added MUST contain a Service Flow ID. The DSA-Response for successfully Admitted or Active upstream QoS Parameter Sets MUST also contain a Service ID.

If the corresponding DSA-Request uses the Service Class Name (refer to C.2.2.3.4) to request service addition, a DSA-Response MUST contain the QoS Parameter Set associated with the named Service Class. If the Service Class Name is used in conjunction with other QoS Parameters in the DSA-Request, the CMTS MUST accept or reject the DSA-Request using the explicit QoS Parameters in the DSA-Request. If these Service Flow Encodings conflict with the Service Class attributes, the CMTS MUST use the DSA-Request values as overrides for those of the Service Class.

If the transaction is successful, the CMTS MUST assign a Classifier Identifier to each requested Classifier and a PHS Index to each requested PHS Rule. The CMTS MUST use the original Classifier Reference(s) and Service Flow Reference(s) to link the successful parameters in the DSA-RSP.

If the transaction is unsuccessful, the CMTS MUST use the original Classifier Reference(s) and Service Flow Reference(s) to identify the failed parameters in the DSA-RSP.

6.3.13.2 CMTS-Initiated Dynamic Service Addition

If the transaction is unsuccessful, the CM MUST use the Classifier Identifier(s) and Service Flow Identifier(s) to identify the failed parameters in the DSA-RSP.

6.3.14 Dynamic Service Addition — Acknowledge (DSA-ACK)

A Dynamic Service Addition Acknowledge **MUST** be generated in response to a received DSA-RSP. The format of a DSA-ACK **MUST** be as shown in Figure 6-33.

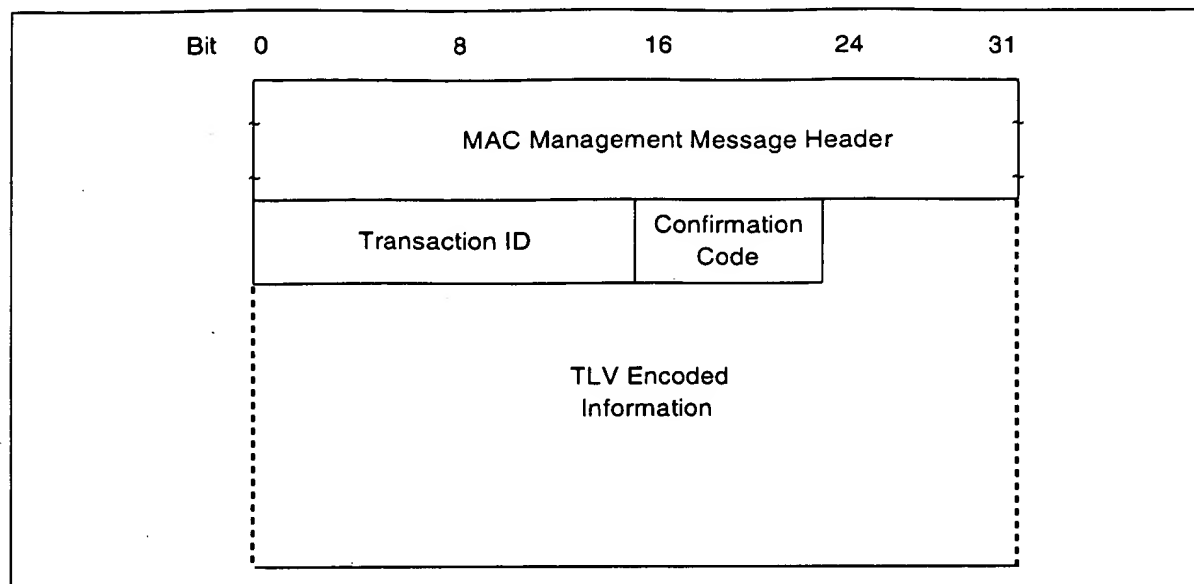


Figure 6-33. Dynamic Service Addition — Acknowledge

Parameters **MUST** be as follows:

Transaction ID	Transaction ID from corresponding DSA-Response.
Confirmation Code	The appropriate Confirmation Code (refer to C.4) for the entire corresponding DSA-Response. ¹

All other parameters are coded TLV tuples.

Service Flow Error Set	The Service Flow Error Set of the DSA-ACK message encodes specifics of any failed Service Flows in the DSA-RSP message. A Service Flow Error Set and identifying Service Flow Reference MUST be included for every failed QoS Parameter of every failed Service Flow in the corresponding DSA-REQ message. This parameter MUST be omitted if the entire DSA-REQ is successful.
-------------------------------	--

If Privacy is enabled, the DSA-ACK message **MUST** contain:

HMAC-Digest	The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute MUST be the final Attribute in the Dynamic Service message's Attribute list. (Refer to Appendix C.1.4.1)
--------------------	---

1. The confirmation code is necessary particularly when a Service Class Name (refer to Section 8.1.3) is used in the DSA-Request. In this case, the DSA-Response could contain Service Flow parameters that the CM is unable to support (either temporarily or as configured).

6.3.15 Dynamic Service Change — Request (DSC-REQ)

A Dynamic Service Change Request MAY be sent by a CM or CMTS to dynamically change the parameters of an existing Service Flow. DSCs changing classifiers MUST carry the entire classifier TLV set for that new classifier.¹

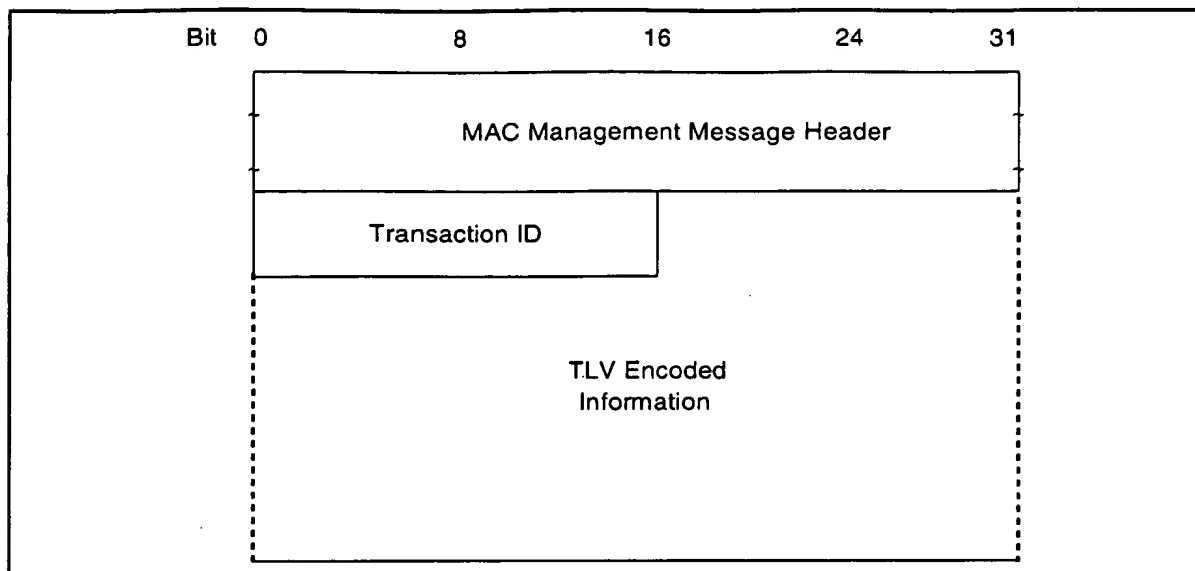


Figure 6-34. Dynamic Service Change — Request

A CM or CMTS MUST generate DSC-REQ messages in the form shown in Figure 6-34 including the following parameters:

Transaction ID Unique identifier for this transaction assigned by the sender.

All other parameters are coded as TLV tuples as defined in Appendix C. A DSC-REQ message MUST NOT carry parameters for more than one Service Flow in each direction, i.e., a DSC-REQ message MUST contain parameters for either a single upstream Service Flow, or for a single downstream Service Flow, or for one upstream and one downstream Service Flow. A DSC-REQ MUST contain at least one of the following:²

Classifier Parameters Specification of the rules to be used to classify packets into a specific service flow — this includes the Dynamic Service Change Action TLV which indicates whether this Classifier should be added, replaced or deleted from the Service Flow (refer to C.2.1.3.7). If included, the Classifier Parameters MUST contain a Classifier Reference/Identifier³ and a Service Flow Identifier.

Service Flow Parameters Specification of the Service Flow's new traffic characteristics and scheduling requirements. The Service Flow Parameters in this message replace all

1. last sentence added 06/21/99 per rfi-n-99043 ew.

2. paragraph edited per rfi-n-99048 06/30/99. ew

3. If the DSC-REQ is CM-initiated and this is a change to an existing Classifier then this is a Classifier Identifier. If the DSC-REQ is CM-initiated and this is a new Classifier then this is a Classifier Reference. Footnote edited 06/21/99 per rfi-n-99043 ew.

parameters currently in use by the Service Flow. If included, the Service Flow Parameters **MUST** contain a Service Flow Identifier.

Payload Header Suppression Parameters

Specification of the rules to be used for Payload Header Suppression to suppress payload headers related to a specific Classifier — this includes the Dynamic Service Change Action TLV which indicates whether this PHS Rule should be added, replaced or deleted from the Service Flow (refer to C.2.1.3.7). If included, the PHS Parameters **MUST** contain a Classifier Reference/Identifier and a Service Flow Identifier.

If Privacy is enabled, a DSC-REQ **MUST** also contain:

HMAC-Digest

The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute **MUST** be the final Attribute in the Dynamic Service message's Attribute list. (Refer to Appendix C.1.4.1)

6.3.16 Dynamic Service Change — Response (DSC-RSP)

A Dynamic Service Change Response **MUST** be generated in response to a received DSC-REQ. The format of a DSC-RSP **MUST** be as shown in Figure 6-35

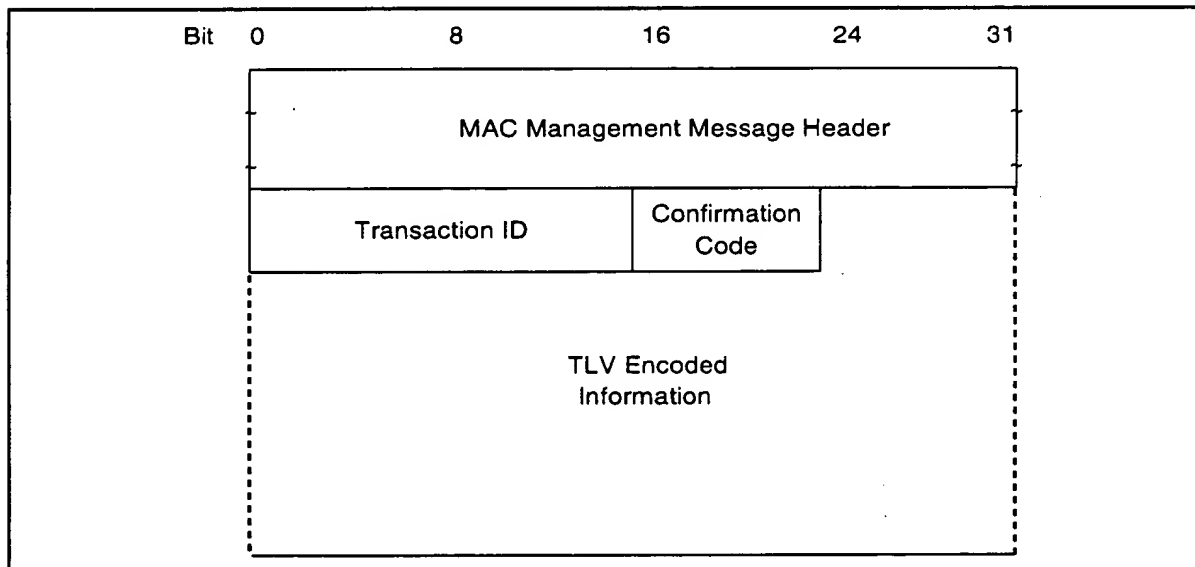


Figure 6-35. Dynamic Service Change — Response

Parameters **MUST** be as follows:

Transaction ID

Transaction ID from corresponding DSC-REQ

Confirmation Code

The appropriate Confirmation Code (refer to C.4) for the corresponding DSC-Request.

All other parameters are coded as TLV tuples as defined in Appendix C.

If the transaction is successful, the DSC-RSP MAY contain one or more of the following:

- Classifier Parameters** The complete specification of the Classifier MUST be included in the DSC-RSP only if it includes a newly assigned Classifier Identifier. If a requested Classifier contained a Classifier Reference, the DSC-RSP MUST contain a Classifier Identifier.
- Service Flow Parameters** The complete specification of the Service Flow MUST be included in the DSC-RSP only if it includes a newly assigned Service Flow Identifier or an expanded Service Class Name. If a Service Flow Parameter set contained an upstream Admitted QoS Parameter Set and this Service Flow does not have an associated SID, the DSC-RSP MUST include a SID. If a Service Flow Parameter set contained a Service Class Name and an Admitted QoS Parameter Set, the DSC-RSP MUST include the QoS Parameter Set corresponding to the named Service Class. If specific QoS Parameters were also included in the classed Service Flow request, these QoS Parameters MUST be included in the DSC-RSP instead of any QoS Parameters of the same type of the named Service Class.
- Payload Header Suppression Parameters** The complete specification of the PHS Parameters MUST be included in the DSC-RSP only if it includes a newly assigned PHS Index. If included, the PHS Parameters MUST contain a Classifier Reference/Identifier and a Service Flow Identifier.

If the transaction is unsuccessful, the DSC-RSP MUST contain at least one of the following:

- Classifier Error Set** A Classifier Error Set and identifying Classifier Reference/Identifier and Service Flow Identifier pair MUST be included for every failed Classifier in the corresponding DSC-REQ. Every Classifier Error Set MUST include every specific failed Classifier Parameter of the corresponding failed Classifier. This parameter MUST be omitted if the entire DSC-REQ is successful.
- Service Flow Error Set** A Service Flow Error Set and identifying Service Flow ID MUST be included for every failed Service Flow in the corresponding DSC-REQ message. Every Service Flow Error Set MUST include every specific failed QoS Parameter of the corresponding Service Flow. This parameter MUST be omitted if the entire DSC-REQ is successful.
- Payload Header Suppression Error Set** A PHS Error Set and identifying PHS Index and Classifier Reference/Identifier pair MUST be included for every failed PHS Rule in the corresponding DSC-REQ. Every PHS Error Set MUST include every specific failed PHS Parameter of the corresponding failed PHS Rule. This parameter MUST be omitted if the entire DSC-REQ is successful.

Regardless of success or failure, if Privacy is enabled for the CM the DSC-RSP MUST contain:

- HMAC-Digest** The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute MUST be the final Attribute in the Dynamic Service message's Attribute list. (Refer to Appendix C.1.4.1)

6.3.17 Dynamic Service Change — Acknowledge (DSC-ACK)

A Dynamic Service Change Acknowledge **MUST** be generated in response to a received DSC-RSP. The format of a DSC-ACK **MUST** be as shown in Figure 6-36

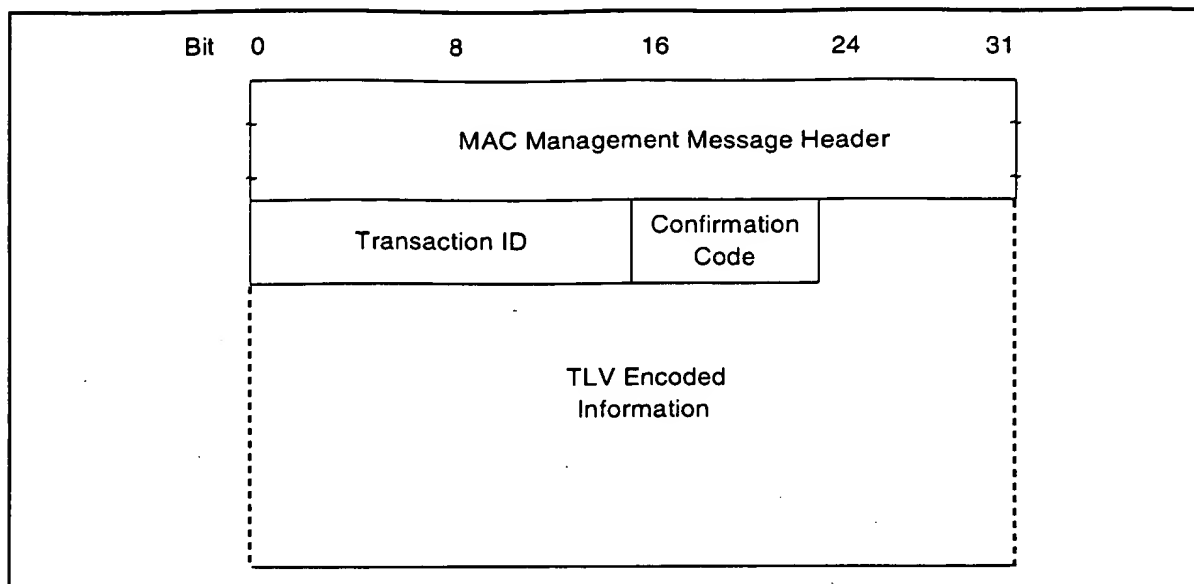


Figure 6-36. Dynamic Service Change — Acknowledge

Parameters **MUST** be as follows:

Transaction ID	Transaction ID from the corresponding DSC-REQ
Confirmation Code	The appropriate Confirmation Code (refer to C.4) for the entire corresponding DSC-Response. ¹

All other parameters are coded TLV tuples.

Service Flow Error Set	The Service Flow Error Set of the DSC-ACK message encodes specifics of any failed Service Flows in the DSC-RSP message. A Service Flow Error Set and identifying Service Flow Identifier MUST be included for every failed QoS Parameter of each failed Service Flow in the corresponding DSC-RSP message. This parameter MUST be omitted if the entire DSC-RSP is successful.
-------------------------------	--

If Privacy is enabled, the DSC-ACK message **MUST** contain:

HMAC-Digest	The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute MUST be the final Attribute in the Dynamic Service message's Attribute list. (Refer to Appendix C.1.4.1)
--------------------	---

1. The Confirmation Code and Service Flow Error Set are necessary particularly when a Service Class Name is (refer to Section 8.1.3) used in the DSC-Request. In this case, the DSC-Response could contain Service Flow parameters that the CM is unable to support (either temporarily or as configured).

6.3.18 Dynamic Service Deletion — Request (DSD-REQ)

A DSD-Request MAY be sent by a CM or CMTS to delete an existing Service Flow. The format of a DSD-Request MUST be as shown in Figure 6-37.




Figure 6-37. Dynamic Service Deletion — Request

Parameters MUST be as follows:

Service Flow Identifier The SFID to be deleted.

Transaction ID Unique identifier for this transaction assigned by the sender.

All other parameters are coded as TLV tuples as defined in Appendix C.

If Privacy is enabled, the DSD-REQ MUST include:

HMAC-Digest The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute MUST be the final Attribute in the Dynamic Service message's Attribute list. (Refer to Appendix C.1.4.1)

6.3.19 Dynamic Service Deletion – Response (DSD-RSP)

A DSD-RSP MUST be generated in response to a received DSD-REQ. The format of a DSD-RSP MUST be as shown in Figure 6-38.

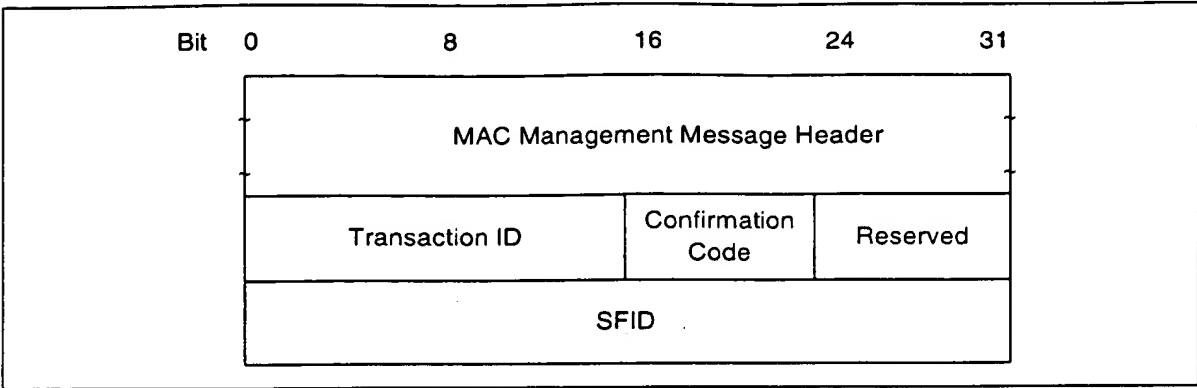


Figure 6-38. Dynamic Service Deletion — Response

Parameters MUST be as follows:

- Service Flow Identifier** SFID from the DSD-REQ to which this acknowledgment refers.
- Transaction ID** Transaction ID from corresponding DSD-REQ.
This page intentionally left blank.
- Confirmation Code** The appropriate Confirmation Code (refer to C.4) for the corresponding DSD-Request.

7 Media Access Control Protocol Operation

7.1 Upstream Bandwidth Allocation

The upstream channel is modeled as a stream of mini-slots. The CMTS **MUST** generate the time reference for identifying these slots. It **MUST** also control access to these slots by the cable modems. For example, it **MAY** grant some number of contiguous slots to a CM for it to transmit a data PDU. The CM **MUST** time its transmission so that the CMTS receives it in the time reference specified. This section describes the elements of protocol used in requesting, granting, and using upstream bandwidth. The basic mechanism for assigning bandwidth management is the allocation MAP. Please refer to Figure 7-1.

The allocation MAP is a MAC Management message transmitted by the CMTS on the downstream channel which describes, for some interval, the uses to which the upstream mini-slots **MUST** be put. A given MAP **MAY** describe some slots as grants for particular stations to transmit data in, other slots as available for contention transmission, and other slots as an opportunity for new stations to join the link.

Many different scheduling algorithms **MAY** be implemented in the CMTS by different vendors; this specification does not mandate a particular algorithm. Instead, it describes the protocol elements by which bandwidth is requested and granted.




Figure 7-1. Allocation Map

The bandwidth allocation **MUST** include the following basic elements:

- Each CM has one or more short (14-bit) service identifiers (SIDs) as well as a 48-bit address.
- Upstream bandwidth is divided into a stream of mini-slots. Each mini-slot is numbered relative to a master reference maintained by the CMTS. The clocking information is distributed to the CMs by means of SYNC packets.
- CMs **MAY** issue requests to the CMTS for upstream bandwidth.

The CMTS MUST transmit allocation MAP PDUs on the downstream channel defining the allowed usage of each mini-slot. The MAP is described below.

7.1.1 The Allocation Map MAC Management Message

The allocation MAP is a varying-length MAC Management message that is transmitted by the CMTS to define transmission opportunities on the upstream channel. It includes a fixed-length header followed by a variable number of information elements (IEs) in the format shown in Section 6.3.4. Each information element defines the allowed usage for a range of mini-slots.

Note that it should be understood by both CM and CMTS that the lower (26-M) bits of alloc start and ack times MUST be used as the effective MAP start and ack times, where M is defined in Section 6.3.3. The relationship between alloc start/ack time counters and the timestamp counter is further described in Section 7.3.4.

7.1.2 Information Elements

Each IE consists of a 14-bit Service ID, a 4-bit type code, and a 14-bit starting offset as defined in Section 6.3.4. Since all stations MUST scan all IEs, it is critical that IEs be short and relatively fixed format. IEs within the MAP are strictly ordered by starting offset. For most purposes, the duration described by the IE is inferred by the difference between the IE's starting offset and that of the following IE. For this reason, a Null IE MUST terminate the list. Refer to Table 6-20.

Four types of Service IDs are defined:

1. 0x3FFF - broadcast, intended for all stations
2. 0x2000-0x3FFE - multicast, purpose is defined administratively. Refer to Appendix A.
3. 0x0001-0x1FFF - unicast, intended for a particular CM or a particular service within that CM
4. 0x0000 - null address, addressed to no station.

All of the Information Elements defined below MUST be supported by conformant CMs. Conformant CMTSs MAY use any of these Information Elements when creating Bandwidth Allocation Maps.

7.1.2.1 The Request IE

The Request IE provides an upstream interval in which requests MAY be made for bandwidth for upstream data transmission. The character of this IE changes depending on the class of Service ID. If broadcast, this is an invitation for CMs to contend for requests. Section 7.4 describes which contention transmit opportunity may be used. If unicast, this is an invitation for a particular CM to request bandwidth. Unicasts MAY be used as part of a Quality of Service scheduling scheme (refer to Section 8.2). Packets transmitted in this interval MUST use the Request MAC Frame format (refer to Section 6.2.5.3).

A small number of Priority Request SIDs are defined in Appendix A. These allow contention for Request IEs to be limited to service flows of a given Traffic Priority (refer to C.2.2.5.2).

7.1.2.2 The Request/Data IE

The Request/Data IE provides an upstream interval in which requests for bandwidth or short data packets MAY be transmitted. This IE is distinguished from the Request IE in that:

- It provides a means by which allocation algorithms MAY provide for "immediate" data contention under light loads, and a means by which this opportunity can be withdrawn as network loading increases.

- Multicast Service IDs **MUST** be used to specify maximum data length, as well as allowed random starting points within the interval. For example, a particular multicast ID **MAY** specify a maximum of 64-byte data packets, with transmit opportunities every fourth slot.

A small number of well-known multicast Service IDs are defined in Appendix A. Others are available for vendor-specific algorithms.

Since data packets transmitted within this interval may collide, the CMTS **MUST** acknowledge any that are successfully received. The data packet **MUST** indicate in the MAC Header that a data acknowledgment is desired (see Table 6-13).

7.1.2.3 The Initial Maintenance IE

The Initial Maintenance IE provides an interval in which new stations may join the network. A long interval, equivalent to the maximum round-trip propagation delay plus the transmission time of the Ranging Request (RNG-REQ) message (see Section 7.3.3), **MUST** be provided to allow new stations to perform initial ranging. Packets transmitted in this interval **MUST** use the RNG-REQ MAC Management message format (refer to Section 6.3.5).

7.1.2.4 The Station Maintenance IE

The Station Maintenance IE provides an interval in which stations are expected to perform some aspect of routine network maintenance, such as ranging or power adjustment. The CMTS **MAY** request that a particular CM perform some task related to network maintenance, such as periodic transmit power adjustment. In this case, the Station Maintenance IE is unicast to provide upstream bandwidth in which to perform this task. Packets transmitted in this interval **MUST** use the RNG-REQ MAC Management message format (see Section 6.3.5).

7.1.2.5 Short and Long Data Grant IEs

The Short and Long Data Grant IEs provide an opportunity for a CM to transmit one or more upstream PDUs. These IEs are issued either in response to a request from a station, or because of an administrative policy providing some amount of bandwidth to a particular station (see class-of-service discussion below). These IEs **MAY** also be used with an inferred length of zero mini slots (a zero length grant), to indicate that a request has been received and is pending (a Data Grant Pending).

Short Data Grants are used with intervals less than or equal to the maximum burst size for this usage specified in the Upstream Channel Descriptor. If Short Data burst profiles are defined in the UCD, then all Long Data Grants **MUST** be for a larger number of mini-slots than the maximum for Short Data. The distinction between Long and Short Data Grants may be exploited in physical-layer forward-error-correction coding; otherwise, it is not meaningful to the bandwidth allocation process.

If this IE is a Data Grant Pending (a zero length grant), it **MUST** follow the NULL IE. This allows cable modems to process all actual allocations first, before scanning the Map for data grants pending and data acknowledgments.

7.1.2.6 Data Acknowledge IE

The Data Acknowledge IE acknowledges that a data PDU was received. The CM MUST have requested this acknowledgment within the data PDU (normally this would be done for PDUs transmitted within a contention interval in order to detect collisions).

This IE MUST follow the NULL IE. This allows cable modems to process all actual interval allocations first, before scanning the Map for data grants pending and data acknowledgments.

7.1.2.7 Expansion IE

The Expansion IE provides for extensibility, if more than 16 code points or 32 bits are needed for future IEs.

7.1.2.8 Null IE

A Null IE terminates all actual allocations in the IE list. It is used to infer a length for the last interval. All Data Acknowledge IEs and All Data Grant Pending IEs (Data Grants with an inferred length of 0) must follow the Null IE.

7.1.3 Requests

Requests refer to the mechanism that CMs use to indicate to the CMTS that it needs upstream bandwidth allocation. A Request MAY come as a stand-alone Request Frame transmission (refer to 6.2.5.3) or it MAY come as a piggyback request in the EHDR of another Frame transmission (refer to 6.2.6).

The Request Frame MAY be transmitted during any of the following intervals:

- Request IE
- Request/Data IE
- Short Data Grant IE
- Long Data Grant IE

A piggyback request MAY be contained in the following Extended Headers:

- Request EH element
- Upstream Privacy EH element
- Upstream Privacy EH element with Fragmentation

The request MUST include:

- The Service ID making the request
- The number of mini-slots requested

The number of mini-slots requested **MUST** be the total number that are desired by the CM at the time of the request (including any physical layer overhead)¹, subject to UCD² and administrative limits³. The CM **MUST** request a number of mini-slots corresponding to one complete frame⁴, except in the case of fragmentation in Piggyback Mode (refer to Section 8.3.2.2).

The CM **MUST** have only one request outstanding at a time per Service ID. If the CMTS does not immediately respond with a Data Grant, the CM is able to unambiguously determine that its request is still pending because the CMTS **MUST** continue to issue a Data Grant Pending in every MAP for as long as a request is unsatisfied.

In MAPs, the CMTS **MUST NOT** make a data grant greater than 255 mini-slots to any assigned Service ID. This puts an upper bound on the grant size the CM has to support.

7.1.4 Information Element Feature Compatibility Summary

The following table summarizes feature compatibility with the different data-related IEs.

Table 7-1. IE Feature Compatibility Summary

Information Element	Request Support	Data PDU Support	Concatenation Support	Fragmentation Support
Request IE	MUST	MUST NOT	MUST NOT	MUST NOT
Request/Data IE	MUST	MAY	MAY	MUST NOT
Short Data Grant IE	MAY	MUST	MUST	MUST
Long Data Grant IE	MAY	MUST	MUST	MUST

Note: A "MUST" in this table indicates that support for this feature in a given IE is mandatory, not that this feature must be used. e.g. while fragmentation support is mandatory, clearly every PDU transmitted in a Short Data Grant will not be fragmented.

7.1.5 Map Transmission and Timing

The allocation MAP **MUST** be transmitted in time to propagate across the physical cable and be received and handled by the receiving CMs. As such, it **MAY** be transmitted considerably earlier than its effective time. The components of the delay are:

- Worst-case round-trip propagation delay — may be network-specific, but on the order of hundreds of microseconds.
- Queuing delays within the CMTS — implementation-specific.
- Processing delays within the CMs — **MUST** allow a minimum processing time by each CM as specified in Appendix B (CM MAP Processing Time).
- PMD-layer FEC interleaving.

Within these constraints, vendors **MAY** wish to minimize this delay so as to minimize latency of access to the upstream channel.

1. Physical layer overhead that **MUST** be accounted for in a request includes: guard band, preamble, and FEC which are dependent on the burst profile.
2. The CM is limited by the Maximum Burst size for the Long Data Grant IUC in the UCD.
3. The CM is limited by the Maximum Concatenated Burst for the Service Flow (refer to Appendix C.2.2.6.1)
4. A frame is a single MAC frame or a concatenated MAC frame.

The number of mini-slots described MAY vary from MAP to MAP. At minimum, a MAP MAY describe a single mini-slot. This would be wasteful in both downstream bandwidth and in processing time within the CMs. At maximum, a MAP MAY stretch to tens of milliseconds. Such a MAP would provide poor upstream latency. Allocation algorithms MAY vary the size of the maps over time to provide a balance of network utilization and latency under varying traffic loads.

At minimum, a MAP MUST contain two Information Elements: one to describe an interval and a null IE to terminate the list. At a maximum, a MAP MUST be bounded by a limit of 240 information elements. Maps are also bounded in that they MUST NOT describe more than 4096 mini-slots into the future. The latter limit is intended to bound the number of future mini-slots that each CM is required to track. Even though multiple maps MAY be outstanding, the sum of the number of mini-slots they describe MUST NOT exceed 4096.

The set of all maps, taken together, MUST describe every mini-slot in the upstream channel. If a CM fails to receive a MAP describing a particular interval, it MUST NOT transmit during that interval.

7.1.6 Protocol Example

This section illustrates the interchange between the CM and the CMTS when the CM has data to transmit (Figure 7-2). Suppose a given CM has a data PDU available for transmission.

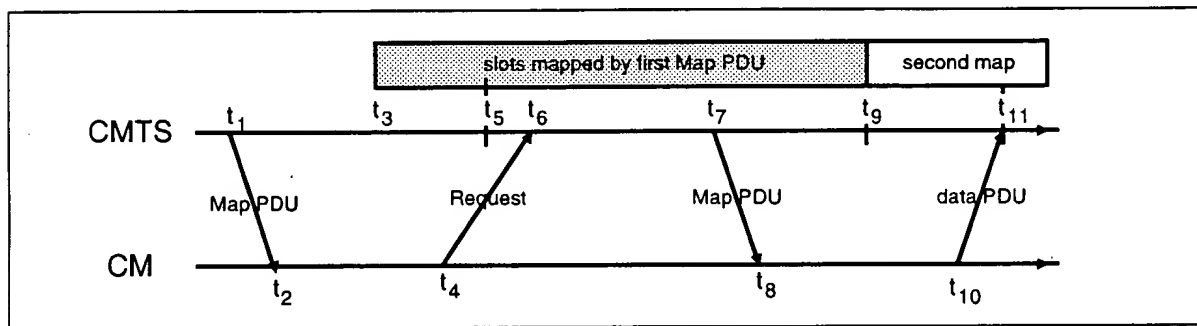


Figure 7-2. Protocol Example

Description

- At time t_1 , the CMTS transmits a MAP whose effective starting time is t_3 . Within this MAP is a Request IE which will start at t_5 . The difference between t_1 and t_3 is needed to allow for:
 - Downstream propagation delay (including FEC interleaving) to allow all CMs to receive the Map
 - Processing time at the CM (allows the CMs to parse the Map and translate it into transmission opportunities)
 - Upstream propagation delay (to allow the CM's transmission of the first upstream data to begin in time to arrive at the CMTS at time t_3).
- At t_2 , the CM receives this MAP and scans it for request opportunities. In order to minimize request collisions, it calculates t_6 as a random offset based on the Data Backoff Start value in the most recent Map (see Section 7.4, also the multicast SID definitions in Section A.2).
- At t_4 , the CM transmits a request for as many mini-slots as needed to accommodate the PDU. Time t_4 is chosen based on the ranging offset (see Section 7.3.3) so that the request will arrive at the CMTS at t_6 .
- At t_6 , the CMTS receives the request and schedules it for service in the next MAP. (The choice of which requests to grant will vary with the class of service requested, any competing requests, and the algorithm used by the CMTS.)

5. At t_7 , the CMTS transmits a MAP whose effective starting time is t_9 . Within this MAP, a data grant for the CM will start at t_{11} .
6. At t_8 , the CM receives the MAP and scans for its data grant.
7. At t_{10} , the CM transmits its data PDU so that it will arrive at the CMTS at t_{11} . Time t_{10} is calculated from the ranging offset as in step 3.

Steps 1 and 2 need not contribute to access latency if CMs routinely maintain a list of request opportunities.

At Step 3, the request may collide with requests from other CMs and be lost. The CMTS does not directly detect the collision. The CM determines that a collision (or other reception failure) occurred when the next MAP fails to include acknowledgment of the request. The CM **MUST** then perform a back-off algorithm and retry. (Refer to Section 7.4.1)

At Step 4, the CMTS scheduler **MAY** fail to accommodate the request within the next MAP. If so, it **MUST** reply with a zero-length grant in that MAP or discard the request by giving no grant at all. It **MUST** continue to report this zero-length grant in all succeeding maps until the request can be granted or is discarded. This **MUST** signal to the CM that the request is still pending. So long as the CM is receiving a zero-length grant, it **MUST NOT** issue new requests for that service queue.

7.2 Support for Multiple Channels

Vendors **MAY** choose to offer various combinations of upstream and downstream channels within one MAC service access point. The upstream bandwidth allocation protocol allows for multiple upstream channels to be managed via one or many downstream channels.

If multiple upstream channels are associated with a single downstream channel, then the CMTS **MUST** send one allocation MAP per upstream channel. The MAP's channel identifier, taken with the Upstream Channel Descriptor Message (see Section 6.3.3), **MUST** specify to which channel each MAP applies. There is no requirement that the maps be synchronized across channels. Appendix H provides an example.

If multiple downstream channels are associated with a single upstream channel, the CMTS **MUST** ensure that the allocation MAP reaches all CMs. That is, if some CMs are attached to a particular downstream channel, then the MAP **MUST** be transmitted on that channel. This **MAY** necessitate that multiple copies of the same MAP be transmitted. The Alloc Start Time in the MAP header **MUST** always relate to the SYNC reference on the downstream channel on which it is transmitted.

If multiple downstream channels are associated with multiple upstream channels, the CMTS **MAY** need to transmit multiple copies of multiple maps to ensure both that all upstream channels are mapped and that all CMs have received their needed maps.

7.3 Timing and Synchronization

One of the major challenges in designing a MAC protocol for a cable network is compensating for the large delays involved. These delays are an order of magnitude larger than the transmission burst time in the upstream. To compensate for these delays, the cable modem **MUST** be able to time its transmissions precisely to arrive at the CMTS at the start of the assigned mini-slot.

To accomplish this, two pieces of information are needed by each cable modem:

- a global timing reference sent downstream from the CMTS to all cable modems.
- a timing offset, calculated during a ranging process, for each cable modem.

7.3.1 Global Timing Reference

The CMTS **MUST** create a global timing reference by transmitting the Time Synchronization (SYNC) MAC management message downstream at a nominal frequency. The message contains a timestamp that exactly identifies when the CMTS transmitted the message. Cable modems **MUST** then compare the actual time the message was received with the timestamp and adjust their local clock references accordingly.

The Transmission Convergence sublayer must operate closely with the MAC sublayer to provide an accurate timestamp for the SYNC message. As mentioned in the Ranging section below (Section 7.3.3), the model assumes that the timing delays through the remainder of the PHY layer **MUST** be relatively constant. Any variation in the PHY delays **MUST** be accounted for in the guard time of the PHY overhead.

It is intended that the nominal interval between SYNC messages be tens of milliseconds. This imposes very little downstream overhead while letting cable modems acquire their global timing synchronization quickly.

7.3.2 CM Channel Acquisition

Any cable modem **MUST NOT** use the upstream channel until it has successfully synchronized to the downstream.

First, the cable modem **MUST** establish PMD sublayer synchronization. This implies that it has locked onto the correct frequency, equalized the downstream channel, recovered any PMD sublayer framing and the FEC is operational (refer to Section 9.2.2). At this point, a valid bit stream is being sent to the transmission convergence sublayer. The transmission convergence sublayer performs its own synchronization (see Section 5). On detecting the well-known DOCSIS PID, along with a payload unit start indicator per [ITU-T H.222.0], it delivers the MAC frame to the MAC sublayer.

The MAC sublayer **MUST** now search for the Timing Synchronization (SYNC) MAC management messages. The cable modem achieves MAC synchronization once it has received at least two SYNC messages and has verified that its clock tolerances are within specified limits.

A cable modem remains in "SYNC" as long as it continues to successfully receive the SYNC messages. If the Lost SYNC Interval (refer to Appendix B) has elapsed without a valid SYNC message, a cable modem **MUST NOT** use the upstream and **MUST** try to re-establish synchronization again.

7.3.3 Ranging

Ranging is the process of acquiring the correct timing offset such that the cable modem's transmissions are aligned to the correct mini-slot boundary. The timing delays through the PHY layer **MUST** be relatively constant. Any variation in the PHY delays **MUST** be accounted for in the guard time of the upstream PMD overhead.

First, a cable modem **MUST** synchronize to the downstream and learn the upstream channel characteristics through the Upstream Channel Descriptor MAC management message. At this point, the cable modem **MUST** scan the Bandwidth Allocation MAP message to find an Initial Maintenance Region. Refer to Section 7.1.2.4. The CMTS **MUST** make an Initial Maintenance region large enough to account for the variation in delays between any two CMs.

The cable modem **MUST** put together a Ranging Request message to be sent in an Initial Maintenance region. The SID field **MUST** be set to the non-initialized CM value (zero).

Ranging adjusts each CM's timing offset such that it appears to be located right next to the CMTS. The CM **MUST** set its initial timing offset to the amount of internal fixed delay equivalent to putting this CM next to the CMTS. This amount includes delays introduced through a particular implementation, and **MUST** include the downstream PHY interleaving latency.

When the Initial Maintenance transmit opportunity occurs, the cable modem **MUST** send the Ranging Request message. Thus, the cable modem sends the message as if it was physically right at the CMTS.

Once the CMTS has successfully received the Ranging Request message, it **MUST** return a Ranging Response message addressed to the individual cable modem. Within the Ranging Response message **MUST** be a temporary SID assigned to this cable modem until it has completed the registration process. The message **MUST** also contain information on RF power level adjustment and offset frequency adjustment as well as any timing offset corrections.

The cable modem **MUST** now wait for an individual Station Maintenance region assigned to its temporary SID. It **MUST** now transmit a Ranging Request message at this time using the temporary SID along with any power level and timing offset corrections.

The CMTS **MUST** return another Ranging Response message to the cable modem with any additional fine tuning required. The ranging request/response steps **MUST** be repeated until the response contains a Ranging Successful notification or the CMTS aborts ranging. Once successfully ranged, the cable modem **MUST** join normal data traffic in the upstream. See Section 9 for complete details on the entire initialization sequence. In particular, state machines and the applicability of retry counts and timer values for the ranging process are defined in Section 9.2.5.

Note: The burst type to use for any transmission is defined by the Interval Usage Code (IUC). Each IUC is mapped to a burst type in the UCD message.

7.3.4 Timing Units and Relationships

The SYNC message conveys a time reference that is measured in 6.25-microsecond ticks. Additional resolution of 6.25/64 microseconds is also present in the SYNC message to allow the CM to track the CMTS clock with a small phase offset. These units were chosen as the greatest-common-divisor of the upstream mini-slot time across various modulations and symbol rates. As this is decoupled from particular upstream channel characteristics, a single SYNC time reference may be used for all upstream channels associated with the downstream channel.

The bandwidth allocation MAP uses time units of "mini-slots." A mini-slot represents the byte-time needed for transmission of a fixed number of bytes. The mini-slot is expected to represent 16 byte-times, although other values could be chosen. The size of the mini-slot, expressed as a multiple of the SYNC time reference, is carried in the Upstream Channel Descriptor. The example in Table 7-2 relates mini-slots to the SYNC time ticks:

Table 7-2. Example Relating Mini-Slots to Time Ticks

Parameter	Example Value
Time tick	6.25 microseconds
Bytes per mini-slot	16 (nominal, when using QPSK modulation)
Symbols/byte	4 (assuming QPSK)
Symbols/second	2,560,000
Mini-slots/second	40,000
Microseconds/mini-slot	25
Ticks/mini-slot	4

Note that the symbols/byte is a characteristic of an individual burst transmission, not of the channel. A mini-slot in this instance could represent either 16 or 32 bytes, depending on the modulation choice.

A “mini-slot” is the unit of granularity for upstream transmission opportunities. There is no implication that any PDU can actually be transmitted in a single mini-slot.

The MAP counts mini-slots in a 32-bit counter that normally counts to $(2^{32} - 1)$ and then wraps back to zero. The least-significant bits (i.e., bit 0 to bit 25-M) of the mini-slot counter MUST match the most-significant bits (i.e., bit 6+M to bit 31) of the SYNC timestamp counter. That is, mini-slot N begins at timestamp reference $(N \cdot T \cdot 64)$, where $T = 2^M$ is the UCD multiplier that defines the mini-slot (i.e., the number of timeticks per minislot). Note: The unused upper bits of the 32-bit mini-slot counter (i.e., bit 26-M to bit 31) are not needed by the CM and MAY be ignored.

Note: The constraint that the UCD multiplier be a power of two has the consequence that the number of bytes per mini-slot must also be a power of two.

7.4 Upstream Transmission and Contention Resolution

The CMTS controls assignments on the upstream channel through the MAP and determines which mini-slots are subject to collisions. The CMTS MAY allow collisions on either Requests or Data PDUs.

This section provides an overview of upstream transmission and contention resolution. For simplicity, it refers to the decisions a CM makes, however, this is just a pedagogical tool. Since a CM can have multiple upstream Service Flows (each with its own SID) it makes these decisions on a per service queue or per SID basis. Refer to Appendix K for a state transition diagram and more detail.

7.4.1 Contention Resolution Overview

The mandatory method of contention resolution which MUST be supported is based on a truncated binary exponential back-off, with the initial back-off window and the maximum back-off window controlled by the CMTS. The values are specified as part of the Bandwidth Allocation Map (MAP) MAC message and represent a power-of-two value. For example, a value of 4 indicates a window between 0 and 15; a value of 10 indicates a window between 0 and 1023.

When a CM has information to send and wants to enter the contention resolution process, it sets its internal back-off window equal to the Data Backoff Start defined in the MAP currently in effect.¹

1. The MAP currently in effect is the MAP whose allocation start time has occurred but which includes IEs that have not occurred.

The CM MUST randomly select a number within its back-off window. This random value indicates the number of contention transmit opportunities which the CM MUST defer before transmitting. A CM MUST only consider contention transmit opportunities for which this transmission would have been eligible. These are defined by either Request IEs or Request/Data IEs in the MAP. Note: Each IE can represent multiple transmission opportunities.

As an example, consider a CM whose initial back-off window is 0 to 15 and it randomly selects the number 11. The CM must defer a total of 11 contention transmission opportunities. If the first available Request IE is for 6 requests, the CM does not use this and has 5 more opportunities to defer. If the next Request IE is for 2 requests, the CM has 3 more to defer. If the third Request IE is for 8 requests, the CM transmits on the fourth request, after deferring for 3 more opportunities.

After a contention transmission, the CM waits for a Data Grant (Data Grant Pending) or Data Acknowledge in a subsequent MAP. Once either is received, the contention resolution is complete. The CM determines that the contention transmission was lost when it finds a MAP without a Data Grant (Data Grant Pending) or Data Acknowledge for it and with an Ack time more recent than the time of transmission.¹ The CM MUST now increase its back-off window by a factor of two, as long as it is less than the maximum back-off window. The CM MUST randomly select a number within its new back-off window and repeat the deferring process described above.

This re-try process continues until the maximum number of retries (16) has been reached, at which time the PDU MUST be discarded. Note: The maximum number of retries is independent of the initial and maximum back-off windows that are defined by the CMTS.

If the CM receives a unicast Request or Data Grant at any time while deferring for this SID, it MUST stop the contention resolution process and use the explicit transmit opportunity.

The CMTS has much flexibility in controlling the contention resolution. At one extreme, the CMTS MAY choose to set up the Data Backoff Start and End to emulate an Ethernet-style back-off with its associated simplicity and distributed nature, but also its fairness and efficiency issues. This would be done by setting Data Backoff Start = 0 and End = 10 in the MAP. At the other end, the CMTS MAY make the Data Backoff Start and End identical and frequently update these values in the MAP so all cable modems are using the same, and hopefully optimal, back-off window.

7.4.2 Transmit Opportunities

A Transmit Opportunity is defined as any mini-slot in which a CM may be allowed to start a transmission. Transmit Opportunities typically apply to contention opportunities and are used to calculate the proper amount to defer in the contention resolution process.

The number of Transmit Opportunities associated with a particular IE in a MAP is dependent on the total size of the region as well as the allowable size of an individual transmission. As an example, assume a REQ IE defines a region of 12 mini-slots. If the UCD defines a REQ Burst Size that fits into a single mini-slot then there are 12 Transmit Opportunities associated with this REQ IE, i.e., one for each mini-slot. If the UCD defines a REQ that fits in two mini-slots, then there are six Transmit Opportunities and a REQ can start on every other mini-slot.

1. Data Acknowledge IEs are intended for collision detection only and is not designed for providing reliable transport (that is the responsibility of higher layers). If a MAP is lost or damaged, a CM waiting for a Data Acknowledge MUST assume that its contention data transmission was successful and MUST NOT retransmit the data packet. This prevents the CM from sending duplicate packets unnecessarily.

As another example, assume a REQ/Data IE that defines a 24 mini-slot region. If it is sent with an SID of 0x3FF4 (refer to Appendix A), then a CM can potentially start a transmit on every fourth mini-slot; so this IE contains a total of six Transmit Opportunities (TX OP). Similarly, a SID of 0x3FF6 implies four TX OPs; 0x3FF8 implies three TX OPs; and 0x3FFC implies two TX OPs.

For an Initial Maintenance IE, a CM **MUST** start its transmission in the first mini-slot of the region; therefore it has a single Transmit Opportunity. The remainder of the region is used to compensate for the round trip delays since the CM has not yet been ranged.

Station Maintenance IEs, Short/Long Data Grant IEs and unicast Request IEs are unicast and thus are not typically associated with contention Transmit Opportunities. They represent a single dedicated, or reservation based, Transmit Opportunity.

In summary:

Table 7-3. Transmit Opportunity

Interval	SID Type	Transmit Opportunity
Request	Broadcast	# minislots required for a Request
Request	Multicast	# minislots required for a Request
Request/Data	Broadcast	Not allowed
Request/Data	Well-known Multicast	As defined by SID in Appendix A
Request/Data	Multicast	Vendor specific algorithms
Initial Maint.	Broadcast	Entire interval is a single tx opp.
Initial Maint.	Multicast	Entire interval is a single tx opp.

7.4.3 CM Bandwidth Utilization

The following rules govern the response a CM makes when processing maps.

Note: These standard behaviors can be overridden by the CM's Request/Transmission Policy (refer to Section C.2.2.6.3):

1. A CM **MUST** first use any Grants assigned to it. Next, the CM **MUST** use any unicast REQ for it. Finally, the CM **MUST** use the next available broadcast/multicast REQ or REQ/Data IEs for which it is eligible.
2. Only one Request may be outstanding at a time for a particular Service ID.
3. If a CM has a Request pending, it **MUST NOT** use intervening contention intervals for that Service ID.

7.5 Data Link Encryption Support

The procedures to support data link encryption are defined in [DOCSIS8]. The interaction between the MAC layer and the security system is limited to the items defined below.

7.5.1 MAC Messages

MAC Management Messages (Section 6.3) **MUST NOT** be encrypted.¹

1. Except for certain cases where such a frame is included in a fragmented concatenated burst on the upstream. (Refer to Section 6.2.7.1)

7.5.2 Framing

The following rules **MUST** be followed when encryption is applied to a data PDU:

- Privacy EH element of [DOCSIS8] **MUST** be in the extended header and **MUST** be the first EH element of the Extended Header field (EHDR).
- Encrypted data are carried as Data PDUs to the Cable MAC transparently.

This page intentionally left blank.

8 Quality of Service & Fragmentation

This specification introduces several new Quality of Service (QoS) related concepts not present in [DOCSIS9]. These include:

- Packet Classification & Flow Identification
- Service Flow QoS Scheduling
- Dynamic Service Establishment
- Fragmentation
- Two-Phase Activation Model

8.1 Theory of Operation

The various DOCSIS protocol mechanisms described in this document can be used to support Quality of Service (QoS) for both upstream and downstream traffic through the CM and the CMTS. This section provides an overview of the QoS protocol mechanisms and their part in providing end-to-end QoS.

The requirements for Quality of Service include:

- A configuration and registration function for pre-configuring CM-based QoS **Service Flows** and traffic parameters.
- A signaling function for dynamically establishing QoS-enabled Service Flows and traffic parameters
- A traffic-shaping and traffic-policing function for Service Flow-based traffic management, performed on traffic arriving from the upper layer service interface and outbound to the RF.
- Utilization of MAC scheduling and traffic parameters for upstream Service Flows.
- Utilization of QoS traffic parameters for downstream Service Flows.
- Classification of packets arriving from the upper layer service interface to a specific active Service Flow.
- Grouping of Service Flow properties into named **Service Classes**, so upper layer entities and external applications (at both the CM and CMTS) can request Service Flows with desired QoS parameters in a globally consistent way.

The principal mechanism for providing enhanced QoS is to classify packets traversing the RF MAC interface into a **Service Flow**. A Service Flow is a unidirectional flow of packets that is provided a particular Quality of Service. The CM and CMTS provide this QoS by shaping, policing, and prioritizing traffic according to the **QoS Parameter Set** defined for the Service Flow.

The primary purpose of the Quality of Service features defined here is to define transmission ordering and scheduling on the Radio Frequency Interface. However, these features often need to work in conjunction with mechanisms beyond the RF interface in order to provide end-to-end QoS or to police the behavior of cable modems. For example, the following behaviors are permitted:

- Policies may be defined by CM MIBs which overwrite the TOS byte. Such policies are outside the scope of the RFI specification. In the upstream direction the CMTS polices the TOS byte setting regardless of how the TOS byte is derived or by whom it is written (originator or CM policy).¹
- The queueing of Service Flow packets at the CMTS in the downstream direction may be based on the TOS byte.

1. Bullet edited 06/22/99 per rfi-n-99043 ew

- Downstream Service Flows can be reclassified by the CM to provide enhanced service onto the subscriber-side network.

Service Flows exist in both the upstream and downstream direction, and MAY exist without actually being activated to carry traffic. Service Flows have a 32-bit **Service Flow Identifier** (SFID) assigned by the CMTS. All Service Flows have an SFID; active upstream Service Flows also have a 14-bit **Service Identifier** (SID).

At least two Service Flows must be defined in each configuration file: one for upstream and one for downstream service. The first upstream Service Flow describes the **Primary Upstream Service Flow**, and is the default Service Flow used for otherwise unclassified traffic and all MAC messages. The first downstream Service Flow describes service to the **Primary Downstream Service Flow**. Additional Service Flows defined in the Configuration file create Service Flows that are provided QoS services.

Conceptually, incoming packets are matched to a **Classifier** that determines to which QoS Service Flow the packet is forwarded. The Classifier can examine the LLC header of the packet, the IP/TCP/UDP header of the packet or some combination of the two. If the packet matches one of the Classifiers, it is forwarded to the Service Flow indicated by the SFID attribute of the Classifier. If the packet is not matched to a Classifier, it is forwarded on the Primary Service Flow.

8.1.1 Concepts

8.1.1.1 Service Flows

A **Service Flow** is a MAC-layer transport service that provides unidirectional transport of packets either to upstream packets transmitted by the CM or to downstream packets transmitted by the CMTS¹. A Service Flow is characterized by a set of **QoS Parameters** such as latency, jitter, and throughput assurances. In order to standardize operation between the CM and CMTS, these attributes include details of how the CM requests upstream minislots and the expected behavior of the CMTS upstream scheduler.

A Service Flow is partially characterized by the following attributes²:

- **ServiceFlowID**: exists for all service flows
- **ServiceID**: only exists for admitted or active upstream service flows
- **ProvisionedQoSParamSet**: defines a set of QoS Parameters which appears in the configuration file and is presented during registration. This MAY define the initial limit for authorizations allowed by the authorization module. The ProvisionedQoSParamSet is defined once when the Service Flow is created via registration.³
- **AdmittedQoSParamSet**: defines a set of QoS parameters for which the CMTS (and possibly the CM) are reserving resources. The principal resource to be reserved is bandwidth, but this also includes any other memory or time-based resource required to subsequently activate the flow.

-
1. A Service Flow, as defined here, has no direct relationship to the concept of a "flow" as defined by the IETF's Integrated Services (intserv) Working Group [RFC-2212]. An intserv flow is a collection of packets sharing transport-layer endpoints. Multiple intserv flows can be served by a single Service Flow. However, the Classifiers for a Service Flow may be based on 802.1P/Q criteria, and so may not involve intserv flows at all.
 2. Some attributes are derived from the above attribute list. The Service Class Name is an attribute of the ProvisionedQoSParamSet. The activation state of the Service Flow is determined by the ActiveQoSParamSet. If the ActiveQoSParamSet is null then the service flow is inactive.
 3. The ProvisionedQoSParamSet is null when a flow is created dynamically.

- **ActiveQoSParamSet**: defines set of QoS parameters defining the service actually being provided to the Service Flow. Only an Active Service Flow may forward packets.

A Service Flow exists when the CMTS assigns a Service Flow ID (SFID) to it. The SFID serves as the principal identifier in the CM and CMTS for the Service Flow. A Service Flow which exists has at least an SFID, and an associated Direction.

The **Authorization Module** is a logical function within the CMTS that approves or denies every change to QoS Parameters and Classifiers associated with a Service Flow. As such it defines an “envelope” that limits the possible values of the AdmittedQoSParameterSet and ActiveQoSParameterSet.

The relationship between the QoS Parameter Sets is as shown in Figure 8-1 and Figure 8-2. The ActiveQoSParameterSet is always a subset¹ of the AdmittedQoSParameterSet which is always a subset of the authorized “envelope.” In the dynamic authorization model, this envelope is determined by the Authorization Module (labeled as the AuthorizedQoSParameterSet). In the provisioned authorization model, this envelope is determined by the ProvisionedQoSParameterSet. (Refer to Section 8.1.4 for further information on the authorization models)

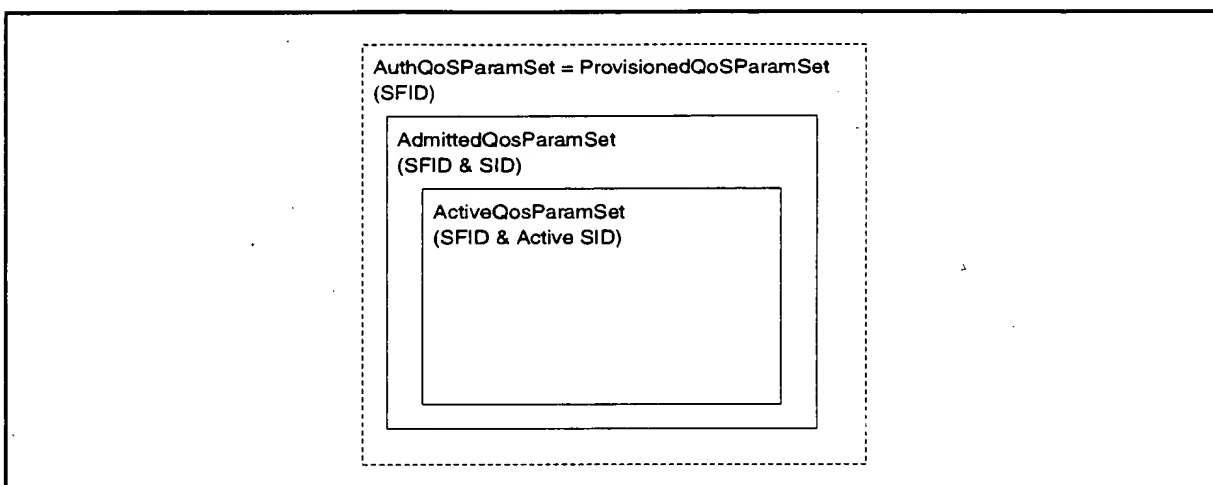


Figure 8-1. Provisioned Authorization Model “Envelopes”

1. To say that QoS Parameter Set A is a subset of QoS Parameter Set B the following **MUST** be true for all QoS Parameters in A and B:
 - if (a smaller QoS parameter value indicates less resources, e.g. Maximum Traffic Rate)
 - A is a subset of B if the parameter in A is less than or equal to the same parameter in B
 - if (a larger QoS parameter value indicates less resources, e.g. Tolerated Grant Jitter)
 - A is a subset of B if the parameter in A is greater than or equal to the same parameter in B
 - if (the QoS parameter specifies a periodic interval, e.g. Nominal Grant Interval),
 - A is a subset of B if the parameter in A is an integer multiple of the same parameter in B
 - if (the QoS parameter is not quantitative, e.g. Service Flow Scheduling Type)
 - A is a subset of B if the parameter in A is equal to the same parameter in B

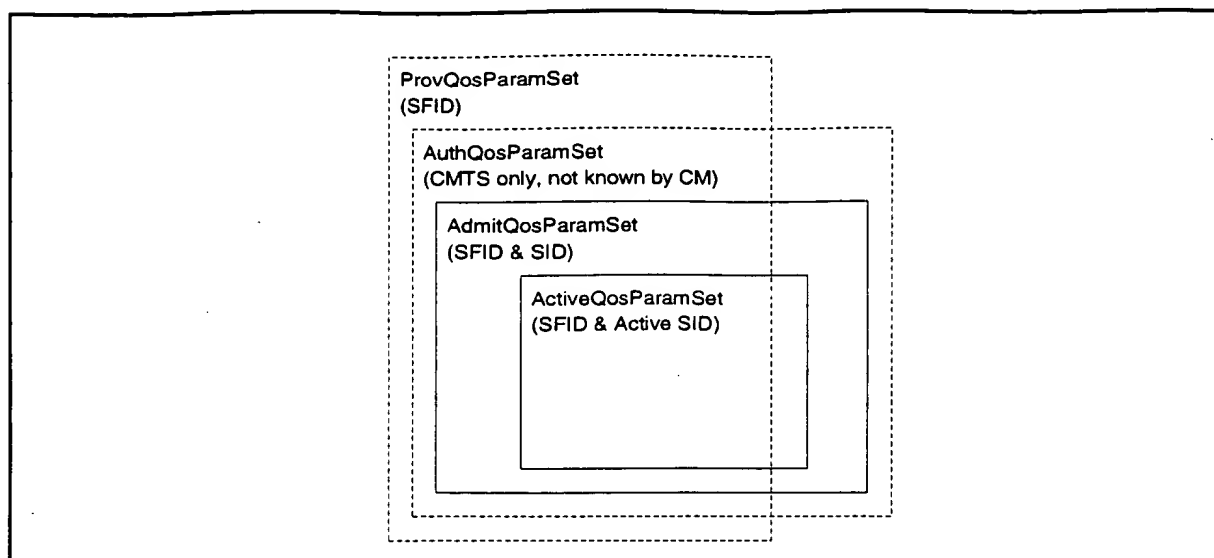


Figure 8-2. Dynamic Authorization Model "Envelopes"

It is useful to think of three types of Service Flows:

- **Provisioned:** this type of Service Flow is known via provisioning through the configuration file, its AdmittedQoSParamSet and ActiveQoSParamSet are both null. A **Provisioned Service Flow** may or may not have associated Classifiers.
- **Admitted:** this type of Service Flow has resources reserved by the CMTS for its AdmittedQoSParamSet, but these parameters are not active (its ActiveQoSParamSet is null). **Admitted Service Flows** may have been provisioned or may have been signalled by some other mechanism. Admitted Service Flows **MUST** have associated Classifiers, though the Classifiers **MUST NOT** yet be active.¹
- **Active:** this type of Service Flow has resources committed by the CMTS for its QoS Parameter Set, (e.g. is actively sending MAPs containing unsolicited grants for a UGS-based service flow). Its ActiveQoSParamSet is non-null. At least one Classifier **MUST** be active for an Active Service Flow.²

1. Bullet edited 06/22/99 per rfi-n-99043 ew

2. Bullet edited 06/22/99 per rfi-n-99043 ew

8.1.1.2 Classifiers

A **Classifier** is a set of matching criteria applied to each packet entering the cable network. It consists of some packet matching criteria (destination IP address, for example), a **classifier priority**, and a reference to a service flow. If a packet matches the specified packet matching criteria, it is then delivered on the referenced service flow.

Several Classifiers may all refer to the same Service Flow. The classifier priority is used for ordering the application of Classifiers to packets. Explicit ordering is necessary because the patterns used by Classifiers may overlap. The priority need not be unique, but care must be taken within a classifier priority to prevent ambiguity in classification. (Refer to Section 8.1.6.1) **Downstream Classifiers** are applied by the CMTS to packets it is transmitting, and **Upstream Classifiers** are applied at the CM and may be applied at the CMTS to police the classification of upstream packets. Figure 8-3 illustrates the mappings discussed above.

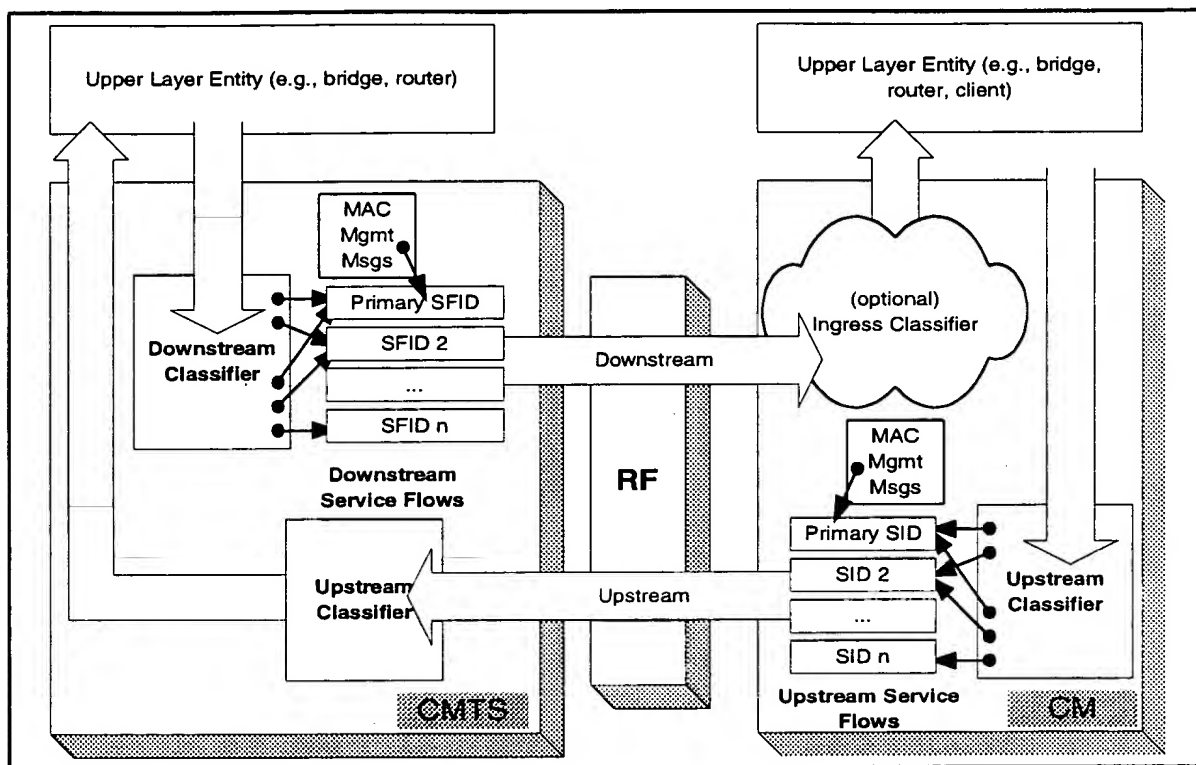


Figure 8-3. Classification within the MAC Layer

CM and CMTS Packet Classification consists of multiple Classifiers. Each Classifier contains a priority field which determines the search order for the Classifier. The highest priority Classifier **MUST** be applied first. If a Classifier is found in which all parameters match the packet, the Classifier **MUST** forward the packet to the corresponding Service Flow. If no Classifier is found in which all parameters match the packet then the packet is classified to the Primary Service Flow.

The packet classification table contains the following fields:

- **Priority** — determines the search order for the table. Higher priority Classifiers are searched before lower priority Classifiers.
- **IP Classification Parameters** — zero or more of the IP classification parameters (IP TOS Range/Mask, IP Protocol, IP Source Address/Mask, IP Destination Address/Mask, TCP/UDP Source Port Start, TCP/UDP Source Port End, TCP/UDP Destination Port Start, TCP/UCP Destination Port End).

- LLC Classification Parameters — zero or more of the LLC classification parameters (Destination MAC Address, Source MAC Address, Ethertype/SAP)
- IEEE 802.1P/Q Parameters — zero or more of the IEEE classification parameters (802.1P Priority Range, 802.1Q VLAN ID)
- Service Flow Identifier — identifier of a specific flow to which this packet is to be directed.

Classifiers can be added to the table either via management operations (configuration file, registration, and SNMP) or via dynamic operations (dynamic signaling, DOCSIS MAC sublayer service interface). SNMP-based operations can view Classifiers that are added via dynamic operations, but can not modify or delete Classifiers that are created by dynamic operations. The format for classification table parameters defined in the configuration file, registration message, or dynamic signaling message is contained in Appendix C.

8.1.2 Object Model

The major objects of the architecture are represented by named rectangles in Figure 8-4. Each object has a number of attributes; the attribute names which uniquely identify the object are underlined. Optional attributes are denoted with brackets. The relationship between the number of objects is marked at each end of the association line between the objects. For example, a Service Flow may be associated with from 0 to 65535 Classifiers, but a Classifier is associated with exactly one Service flow.

The Service Flow is the central concept of the MAC protocol. It is uniquely identified by a 32-bit Service Flow ID (SFID) assigned by the CMTS. Service Flows may be in either the upstream or downstream direction. Admitted Upstream Service Flows are assigned a 14-bit Service ID (SID).

Typically, an outgoing user data Packet is submitted by an upper layer protocol (such as the forwarding bridge of a CM) for transmission on the Cable MAC interface. The packet is compared against a set of Classifiers. The matching Classifier for the Packet identifies the corresponding Service Flow via the Service Flow ID (SFID). In the case where more than one Classifier matches the packet, the highest Priority Classifier is chosen.

The Classifier matching a packet MAY be associated with a Payload Header Suppression Rule. A PHS Rule provides details on how header bytes of a Packet PDU can be omitted, replaced with a Payload Header Suppression Index for transmission and subsequently regenerated at the receiving end. PHS Rules are indexed by the combination of {SFID, PHSI} (refer to Section 8.4). When a Service Flow is deleted, all Classifiers and any associated PHS Rules referencing it MUST also be deleted.

The Service Class is an optional object that may be implemented at the CMTS. It is referenced by an ASCII name which is intended for provisioning purposes. A Service Class is defined in the CMTS to have a particular QoS Parameter Set. The QoS Parameter Sets of a Service Flow may contain a reference to the Service Class Name as a "macro" that selects all of the QoS parameters of the Service Class. The Service Flow QoS Parameter Sets may augment and even override the QoS parameter settings of the Service Class, subject to authorization by the CMTS. (Refer to Appendix C.2.2.5)

If a Packet has already been determined by upper layer policy mechanisms to be associated with a particular Service Class Name/Priority combination, that combination associates the packet with a particular Service Flow directly (refer to Section 8.1.6.1). The upper layer may also be aware of the particular Service Flows in the MAC Sublayer, and may have assigned the Packet directly to a Service Flow. In these cases, a user data Packet is considered to be directly associated with a Service Flow as selected by the upper layer. This is depicted with the dashed arrow in Figure 8-4. (Refer to Appendix E)

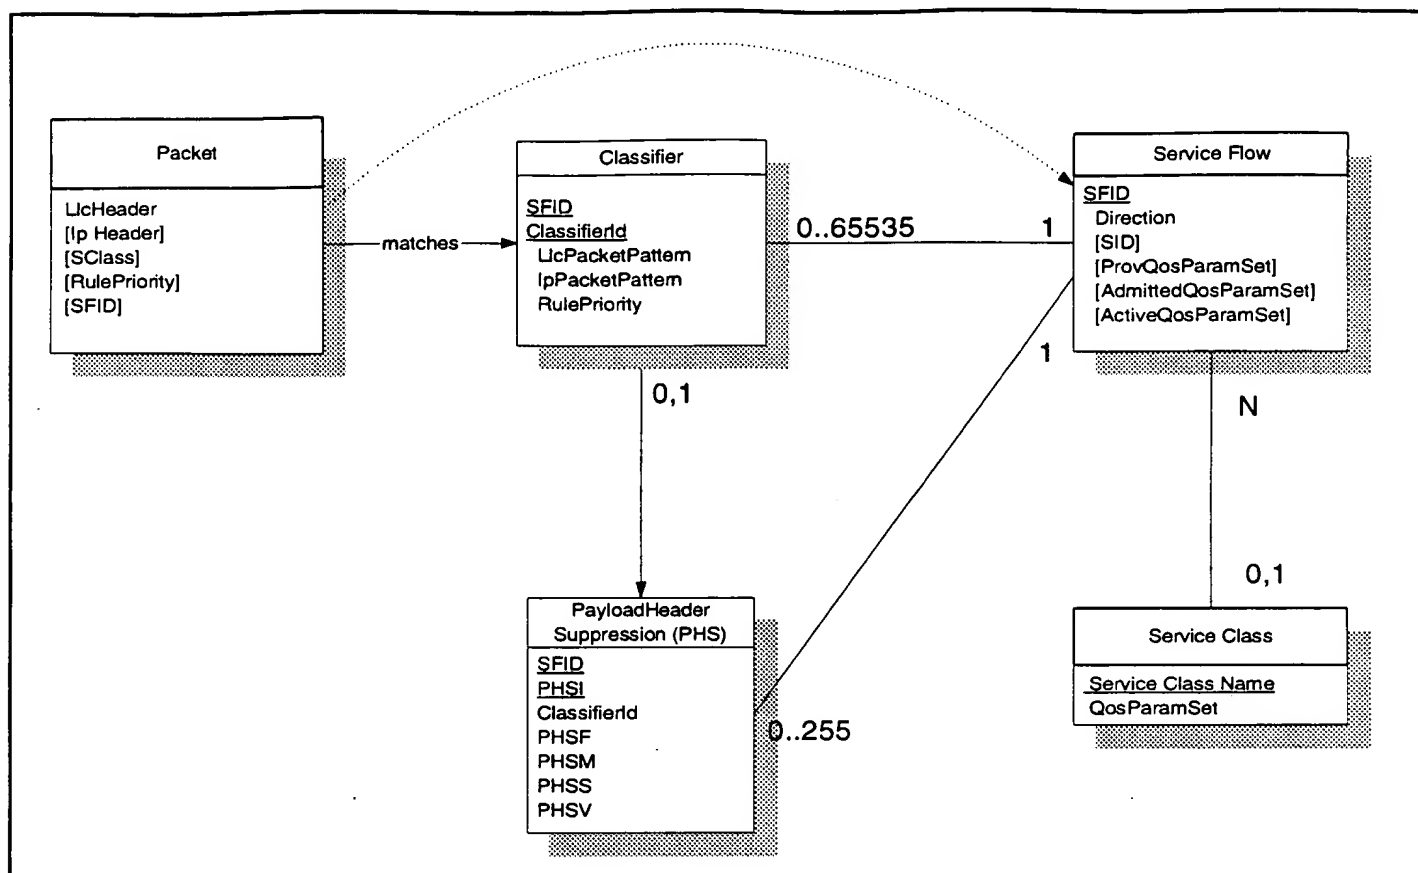


Figure 8-4. Theory of Operation Object Model

8.1.3 Service Classes

The QoS attributes of a Service Flow may be specified in two ways: either by explicitly defining all attributes, or implicitly by specifying a **Service Class Name**. A **Service Class Name** is a string which the CMTS associates with a QoS Parameter Set. It is described further below.

The Service Class serves the following purposes:

1. It allows operators, who so wish, to move the burden of configuring service flows from the provisioning server to the CMTS. Operators provision the modems with the Service Class Name; the implementation of the name is configured at the CMTS. This allows operators to modify the implementation of a given service to local circumstances without changing modem provisioning. For example, some scheduling parameters may need to be tweaked differently for two different CMTSs to provide the same service. As another example, service profiles could be changed by time of day.
2. It allows CMTS vendors to provide class-based-queuing if they choose, where service flows compete within their class and classes compete with each other for bandwidth.
3. It allows higher-layer protocols to create a Service Flow by its Service Class Name. For example, telephony signaling may direct the CM to instantiate any available Provisioned Service Flow of class "G711".
4. It allows packet classification policies to be defined which refer to a desired service class, without having to refer to a particular service flow instance of that class.

Note: The Service Class is optional: the flow scheduling specification may always be provided in full; a service flow may belong to no service class whatsoever. CMTS implementations MAY treat such "unclassified" flows differently from "classified" flows with equivalent parameters.

Any Service Flow MAY have its QoS Parameter Set specified in any of three ways:

- By explicitly including all traffic parameters.
- By indirectly referring to a set of traffic parameters by specifying a Service Class Name.
- By specifying a Service Class Name along with modifying parameters.

The Service Class Name is "expanded" to its defined set of parameters at the time the CMTS successfully admits the Service Flow. The Service Class expansion can be contained in the following CMTS-originated messages: Registration Response, DSA-REQ, DSC-REQ, DSA-RSP and DSC-RSP. In all of these cases, the CMTS MUST include a Service Flow Encoding that includes the Service Class Name and the QoS Parameter Set of the Service Class. If a CM-initiated request contained any supplemental or overriding Service Flow parameters, a successful response MUST also include these parameters.

When a Service Class name is given in an admission or activation request, it is possible that the returned QoS Parameter Set MAY change from activation to activation. This can happen because of administrative changes to the Service Class' QoS Parameter Set at the CMTS. If the definition of a Service Class Name is changed at the CMTS (e.g. its associated QoS Parameter Set is modified), it has no effect on the QoS Parameters of existing Service Flows associated with that Service Class. A CMTS MAY initiate DSC transactions to existing Service Flows which reference the Service Class Name to affect the changed Service Class definition.

When a CM uses the Service Class Name to specify the Admitted QoS Parameter Set, the expanded set of TLV encodings of the Service Flow will be returned to the CM in the response message (REG-RSP, DSA-RSP, or DSC-RSP). Use of the Service Class Name later in the activation request MAY fail if the definition of the Service Class Name has changed and the new required resources are not available. Thus, the CM SHOULD explicitly request the expanded set of TLVs from the response message in its later activation request.

8.1.4 Authorization

Every change to the Service Flow QoS Parameters MUST be approved by an authorization module. This includes every REG-REQ or DSA-REQ message to create a new Service Flow, and every DSC-REQ message to change a QoS Parameter Set of an existing Service Flow. Such changes include requesting an admission control decision (e.g. setting the AdmittedQoSParamSet) and requesting activation of a Service Flow (e.g. setting the ActiveQoSParameterSet). Reduction requests regarding the resources to be admitted or activated are also checked by the authorization module, as are requests to add or change the Classifiers.

In the static authorization model, the authorization module receives all registration messages, and stores the provisioned status of all "deferred" Service Flows. Admission and activation requests for these provisioned service flows will be permitted, as long as the Admitted QoS Parameter Set is a subset of the Provisioned QoS Parameter Set, and the Active QoS Parameter Set is a subset of the Admitted QoS Parameter Set. Requests to change the Provisioned QoS Parameter Set will be refused, as will requests to create new dynamic Service Flows. This defines a static system where all possible services are defined in the initial configuration of each CM.

In the dynamic authorization model, the authorization module not only receives all registration messages, but also communicates through a separate interface to an independent policy server. This policy server may provide to the authorization module advance notice of upcoming admission and activation requests, and specifies the proper authorization action to be taken on those requests. Admission and activation requests from a CM are then checked by the Authorization Module to ensure that the ActiveQoSParameterSet being requested is a subset of the set provided by the policy server. Admission and activation requests from a CM that are signalled in advance

by the external policy server are permitted. Admission and activation requests from a CM that are not pre-signalled by the external policy server may result in a real-time query to the policy server, or may be refused.

During registration, the CM **MUST** send to the CMTS the authenticated set of TLVs derived from its configuration file which defines the Provisioned QoS Parameter Set. Upon receipt and verification at the CMTS, these are handed to the Authorization Module within the CMTS. The CMTS **MUST** be capable of caching the Provisioned QoS Parameter Set, and **MUST** be able to use this information to authorize dynamic flows which are a subset of the Provisioned QoS Parameter Set. The CMTS **SHOULD** implement mechanisms for overriding this automated approval process (such as described in the dynamic authorization model). For example:

- Deny all requests whether or not they have been pre-provisioned
- Define an internal table with a richer policy mechanism but seeded by the configuration file information
- Refer all requests to an external policy server

8.1.5 Types of Service Flows

It useful to think about three basic types of Service Flows. This section describes these three types of Service Flows in more detail. However, it is important to note that there are more than just these three basic types. (Refer to Appendix C.2.2.5.1)

8.1.5.1 Provisioned Service Flows

A Service Flow may be Provisioned but not immediately activated (sometimes called “deferred”). That is, the description of any such service flow in the TFTP configuration file contains an attribute which provisions but defers activation and admission (refer to Appendix C.2.2.5.1). During Registration, the CMTS assigns a Service Flow ID for such a service flow but does not reserve resources. The CMTS **MAY** also require an exchange with a policy module prior to admission.

As a result of external action beyond the scope of this specification (e.g. [PKTCBL-MGCP]), the CM **MAY** choose to activate a Provisioned Service Flow by passing the Service Flow ID and the associated QoS Parameter Sets. The CM **MUST** also provide any applicable Classifiers. If authorized and resources are available, the CMTS **MUST** respond by assigning a SID for the upstream Service Flow. The CMTS **MAY** deactivate the Service Flow, but **SHOULD NOT** delete the Service Flow during the CM registration epoch.

As a result of external action beyond the scope of this specification (e.g. [PKTCBL-MGCP]), the CMTS **MAY** choose to activate a Service Flow by passing the Service Flow ID as well as the SID and the associated QoS Parameter Sets. The CMTS **MUST** also provide any applicable Classifiers. The CMTS **MAY** deactivate the Service Flow, but **SHOULD NOT** delete the Service Flow during the CM registration epoch. Such a Provisioned Service Flow **MAY** be activated and deactivated many times (through DSC exchanges). In all cases, the original Service Flow ID **MUST** be used when reactivating the service flow.

8.1.5.2 Admitted Service Flows

This protocol supports a two-phase activation model which is often utilized in telephony applications. In the two-phase activation model, the resources for a “call” are first “admitted,” and then once the end-to-end negotiation is completed (e.g. called party’s gateway generates an “off-hook” event) the resources are “activated.” Such a two-phase model serves the purposes a) of conserving network resources until a complete end-to-end connection has been established, b) performing policy checks and admission control on resources as quickly as possible, and, in particular, before informing the far end of a connection request, and c) preventing several potential theft-of-service scenarios.

For example, if an upper layer service were using unsolicited grant service, and the addition of upper-layer flows could be adequately provided by increasing the Grants Per Interval QoS parameter, then the following might be used. When the first upper-layer flow is pending, the CM issues a DSA-Request with the Admit Grants Per Interval parameter equal one, and the Activate Grants Per Interval parameter equal zero. Later when the upper-layer flow becomes active, it issues a DSC-Request with the instance of the Activate Grants-per-Interval parameter equal to one. Admission control was performed at the time of the reservation, so the later DSC-Request, having the Activate parameters within the range of the previous reservation, is guaranteed to succeed. Subsequent upper-layer flows would be handled in the same way. If there were three upper-layer flows establishing connections, with one flow already active, the Service Flow would have Admit(ted) Grants-per-Interval equal four, and Active Grants-per-Interval equal one.

An activation request of a Service Flow where the new ActiveQoSParamSet is a subset of the AdmittedQoS-ParamSet and no new classifiers are being added **MUST** be allowed (except in the case of catastrophic failure). An admission request where the AdmittedQoSParamSet is a subset of the previous AdmittedQoSParamSet, so long as the ActiveQoSParamSet remains a subset of the AdmittedQoSParameterSet, **MUST** succeed.

A Service Flow that has resources assigned to its AdmittedQoSParamSet, but whose resources are not yet completely activated, is in a transient state. A timeout value **MUST** be enforced by the CMTS that requires Service Flow activation within this period. (Refer to Appendix C.2.2.5.8) If Service Flow activation is not completed within this interval, the assigned resources in excess of the active QoS parameters **MUST** be released by the CMTS.

It is possible in some applications that a long-term reservation of resources is necessary or desirable. For example, placing a telephone call on hold should allow any resources in use for the call to be temporarily allocated to other purposes, but these resources must be available for resumption of the call later. The AdmittedQoSParamSet is maintained as "soft state" in the CMTS; this state **MUST** be refreshed periodically for it to be maintained without the above timeout releasing the non-activated resources. This refresh **MAY** be signalled with a periodic DSC-REQ message with identical QoS Parameter Sets, or **MAY** be signalled by some internal mechanism within the CMTS outside of the scope of this specification (e.g. by the CMTS monitoring RSVP refresh messages).

8.1.5.3 Active Service Flows

A Service Flow that has a non-NULL set of ActiveQoSParameters is said to be an Active Service Flow. It is requesting¹ and being granted bandwidth for transport of data packets. An admitted Service Flow may be made active by providing an ActiveQoSParameterSet, signaling the resources actually desired at the current time.² This completes the second stage of the two-phase activation model. (Refer to Section 8.1.5.2)

A Service Flow may be Provisioned and immediately activated. This is the case for the Primary Service Flows. It is also typical of Service Flows for monthly subscription services, etc. These Service Flows are established at registration time and **MUST** be authorized by the CMTS MIC. These Service Flows **MAY** also be authorized by the CMTS authorization module.

Alternatively, a Service Flow may be created dynamically and immediately activated. In this case, two-phase activation is skipped and the Service Flow is available for immediate use upon authorization.

1. According to its Request/Transmission Policy (refer to C.2.2.6.3)

2. edited 06/22/99 per rfi-n-99043 ew

8.1.6 Service Flows and Classifiers

The basic model is that the Classifiers associate packets into exactly one Service Flow. The Service Flow Encodings provide the QoS Parameters for treatment of those packets on the RF interface. These encodings are described in Appendix C.2.

In the upstream direction, the CM **MUST** classify upstream packets to Active Service Flows. The CMTS **MUST** classify downstream traffic to Active Downstream Service Flows. There **MUST** be a default downstream service flow for otherwise unclassified broadcast and multicast traffic.

The CMTS polices packets in upstream Service Flows to ensure the integrity of the QoS Parameters and the packet's TOS value.¹ When the rate at which packets are sent is greater than the policed rate at the CMTS, then these packets **MAY** be dropped by the CMTS (refer to C.2.2.5.3). When the value of the TOS byte is incorrect, the CMTS (based on policy) **MUST** police the stream by overwriting the TOS byte (refer to C.2.2.6.10).

It may not be possible for the CM to forward certain upstream packets on certain Service Flows. In particular, a Service Flow using unsolicited grant service with fragmentation disabled cannot be used to forward packets larger than the grant size. If a packet is classified to a Service Flow on which it cannot be transmitted, the CM **MUST** either transmit the packet on the Primary Service Flow or discard the packet depending on the Request/Transmission Policy of the Service Flow to which the packet was classified.

MAC Management messages are not subject to classification and are not part of any service flow. Although MAC Management messages are transferred on the Primary SID, they **MUST** be excluded from any QoS calculations of the Primary Service Flow. Delivery of MAC Management messages is implicitly influenced by the attributes of the associated service flow.

8.1.6.1 Policy-Based Classification and Service Classes

As noted in Appendix E, there are a variety of ways in which packets may be enqueued for transmission at the MAC Service Interface. At one extreme are embedded applications that are tightly bound to a particular Payload Header Suppression Rule (refer to Section 8.4) and which forego more general classification by the MAC. At the other extreme are general transit packets of which nothing is known until they are parsed by the MAC Classification rules. Another useful category is traffic to which policies are applied by a higher-layer entity and then passed to the MAC for further classification to a particular service flow.

Policy-based classification is, in general, beyond the scope of this specification. One example might be the docsDevFilterIpPolicyTable defined in the Cable Device MIB [ID-CDMIB]. Such policies may tend to be longer-lived than individual service flows and MAC classifiers and so it is appropriate to layer the two mechanisms, with a well-defined interface between policies and MAC Service Flow Classification.

The interface between the two layers is the addition of two parameters at the MAC transmission request interface. The two parameters are a Service Class Name and a Rule Priority that is applied to matching the service class name. The Policy Priority is from the same number space as the Packet Classifier Priority of the packet-matching rules used by MAC classifiers. The MAC Classification algorithm is now:

1. sentence edited 06/22/99 per rfi-n-99043 ew

```

MAC_DATA.request( PDU,
                  ServiceClassName,
                  RulePriority)

TxServiceFlowID = FIND_FIRST_SERVICE_FLOW_ID (ServiceClassName)
SearchID = SEARCH_CLASSIFIER_TABLE (All Priority Levels)
IF (SearchID not NULL and Classifier.RulePriority >= MAC_DATA.RulePriority)
    TxServiceFlowID = SearchID

IF (TxServiceFlowID = NULL)
    TRANSMIT_PDU (PrimaryServiceFlowID)
ELSE
    TRANSMIT_PDU (TxServiceFlowID)

```

While Policy Priority competes with Packet Classifier Priority and its choice might in theory be problematic, it is anticipated that well-known ranges of priorities will be chosen to avoid ambiguity. In particular, dynamically-added classifiers **MUST** use the priority range 64-191. Classifiers created as part of registration, as well as policy-based classifiers, **MAY** use zero through 255, but **SHOULD** avoid the dynamic range.

Note: Classification within the MAC sublayer is intended to simply associate a packet with a service flow. If a packet is intended to be dropped it **MUST** be dropped by the higher-layer entity and not delivered to the MAC sublayer.

8.1.7 General Operation

8.1.7.1 Static Operation

Static configuration of Classifiers and Service Flows uses the Registration process. A provisioning server provides the CM with configuration information. The CM passes this information to the CMTS in a Registration Request. The CMTS adds information and replies with a Registration Response. The CM sends a Registration Acknowledge to complete registration.

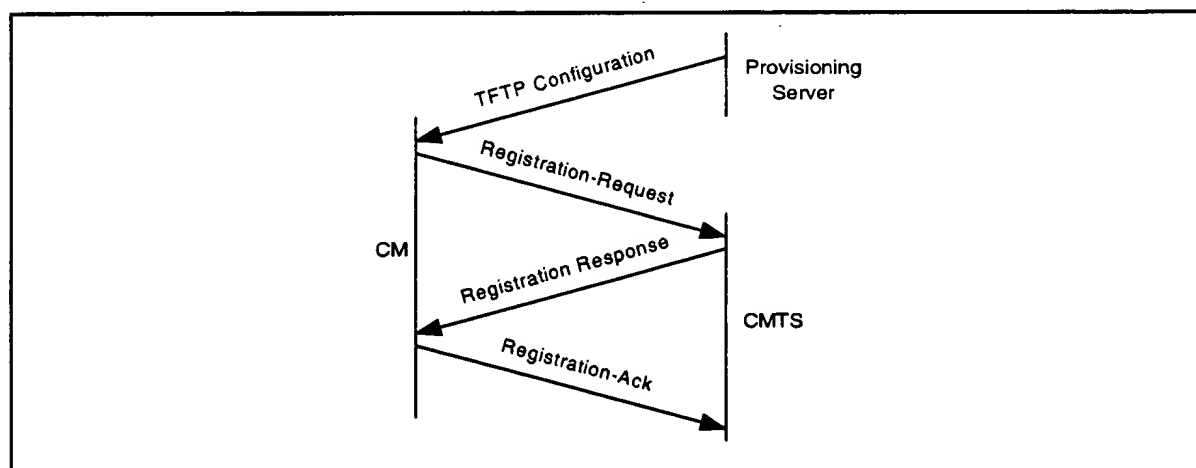


Figure 8-5. Registration Message Flow

A TFTP configuration file consists of one or more instances of Classifiers and Service Flow Encodings. Classifiers are loosely ordered by 'priority'. Each Classifier refers to a Service Flow via a 'service flow reference'. Several Classifiers may refer to the same Service Flow. Additionally, more than one Classifier may have the same priority, and in this case, the particular classifier used is not defined.

Table 8-1. TFTP File Contents

Items	Point To Service Flow Reference	Service Flow Reference	Service Flow ID
Upstream Classifiers Each containing a Service Flow Reference (pointer)	1..n		
Downstream Classifiers Each containing a Service Flow Reference (pointer)	(n+1)..q		
Service Flow Encodings Immediate activation requested, upstream		1..m	None Yet
Service Flow Encodings Provisioned for later activation requested, upstream		(m+1)..n	None Yet
Service Flow Encodings Immediate activation requested, downstream		(n+1)..p	None Yet
Service Flow Encodings Provisioned for later activation requested, downstream		(p+1)..q	None Yet

Service Flow Encodings contain either a full definition of service attributes (omitting defaultable items if desired) or a service class name. A service class name is an ASCII string which is known at the CMTS and which indirectly specifies a set of QoS Parameters. (Refer to Section 8.1.3 and C.2.2.3.4)

Note: At the time of the TFTP configuration file, Service Flow References exist as defined by the provisioning server. Service Flow Identifiers do not yet exist because the CMTS is unaware of these service flow definitions.

The Registration Request packet contains Downstream Classifiers (if to be immediately activated) and all Inactive Service Flows. The configuration file, and thus, the Registration Request, generally does not contain a Downstream Classifier if the corresponding Service Flow is requested with deferred activation. This allows for late binding of the Classifier when the Flow is activated.

Table 8-2. Registration Request Contents

Items	Point To Service Flow Reference	Service Flow Reference	Service Flow ID
Upstream Classifiers Each containing a Service Flow Reference (pointer)	1..n		
Downstream Classifiers Each containing a Service Flow Reference (pointer)	(n+1)..p		
Service Flow Encodings Immediate activation requested, upstream May specify explicit attributes or service class name		1..m	None Yet
Service Flow Encodings Provisioned for later activation requested, upstream Explicit attributes or service class name		(m+1)..n	None Yet
Service Flow Encodings Immediate activation requested, downstream Explicit attributes or service name		(n+1)..p	None Yet
Service Flow Encodings Provisioned for later activation requested, downstream Explicit attributes or service name		(p+1)..q	None Yet

The Registration Response sets the QoS Parameter Sets according to the Quality of Service Parameter Set Type in the Registration Request.

The Registration Response preserves the Service Flow Reference attribute, so that the Service Flow Reference can be associated with SFID and/or SID.

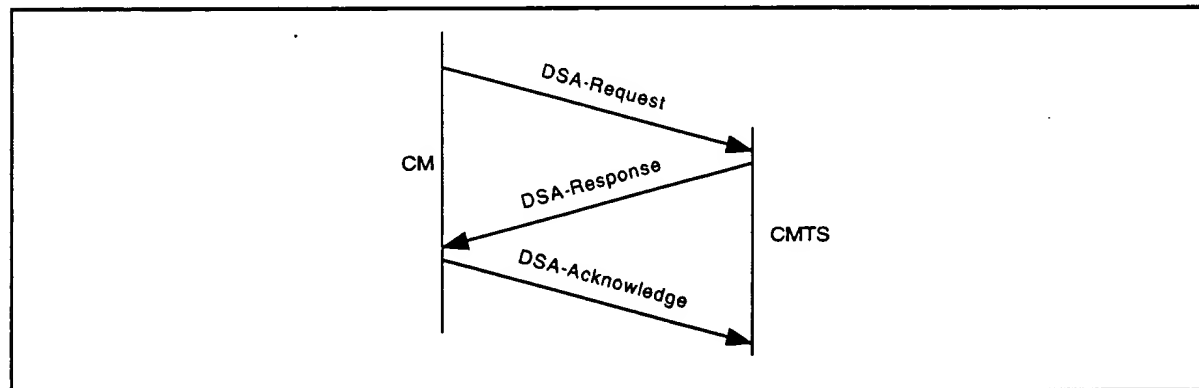
Table 8-3. Registration Response Contents

Items	Service Flow Reference	Service Flow Identifier	Service Identifier
Active Upstream Service Flows Explicit attributes	1..m	SFID	SID
Provisioned Upstream Service Flows Explicit attributes	(m+1)..n	SFID	Not Yet
Active Downstream Service Flows Explicit attributes	(n+1)..p	SFID	N/A
Provisioned Downstream Service Flows Explicit attributes	(p+1)..q	SFID	N/A

The SFID is chosen by the CMTS to identify a downstream or upstream service Flow that has been authorized but not activated. A DSC-Request from a modem to admit or activate a Provisioned Service Flow contains its SFID. If it is a downstream Flow then the Downstream Classifier is also included.

8.1.7.2 Dynamic Service Flow Creation — CM Initiated

Service Flows may be created by the Dynamic Service Addition process, as well as through the Registration process outlined above. The Dynamic Service Addition may be initiated by either the CM or the CMTS, and may create one upstream and/or one downstream dynamic Service Flow(s).¹ A three-way handshake is used to create Service Flows. The CM-initiated protocol is illustrated in Figure 8-6 and described in detail in Section 9.4.1.1.

**Figure 8-6. Dynamic Service Addition Message Flow — CM Initiated**

A DSA-Request from a CM contains Service Flow Reference(s), QoS Parameter set(s) (marked either for admission-only or for admission and activation) and any required Classifiers.²

1. Sentence edited per rfi-n-99048 06/30/99. ew

2. Tables (8-4, 8-5) deleted per rfi-n-99048 06/30/99.ew

8.1.7.3 Dynamic Service Flow Creation — CMTS Initiated

A DSA-Request from a CMTS contains Service Flow Identifier(s) for one upstream and/or one downstream Service Flow, possibly a SID, set(s) of active or admitted QoS Parameters, and any required Classifier(s).¹ The protocol is as illustrated in Figure 8-7 and is described in detail in Section 9.4.1.2.

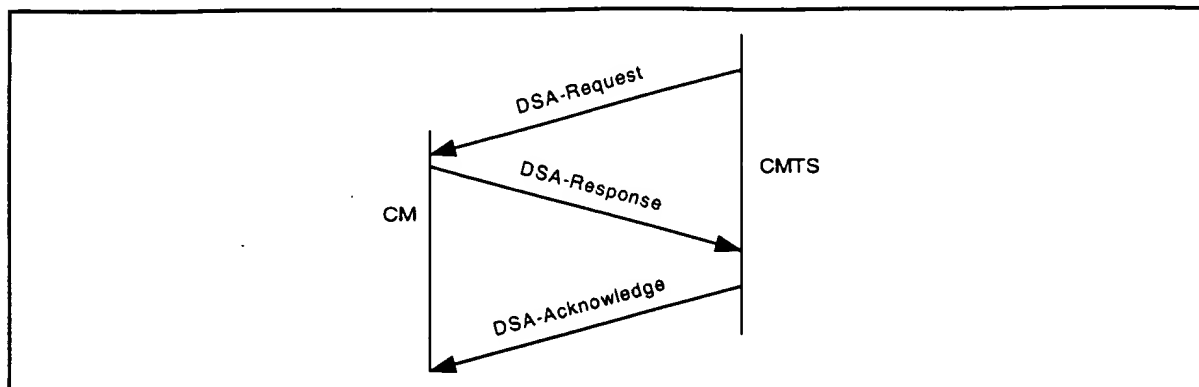


Figure 8-7. Dynamic Service Addition Message Flow — CMTS Initiated

8.1.7.4 Dynamic Service Flow Modification and Deletion

In addition to the methods presented above for creating service flows, protocols are defined for modifying and deleting service flows. Refer to Section 9.4.3 and Section 9.4.2.²

1. Sentence edited per rfi-n-99048 06/30/99. ew

2. Table (8-6) deleted per rfi-n-99048 06/30/99.ew

8.2 Upstream Service Flow Scheduling Services

The following sections define these upstream Service Flow scheduling services: Unsolicited Grant Service (UGS), Real-Time Polling Service (rtPS), Unsolicited Grant Service with Activity Detection (UGS-AD), Non-Real-Time Polling Service (nrtPS), Best Effort (BE) service, and a Committed Information Rate (CIR) service.

8.2.1 Unsolicited Grant Service

The intent of UGS is to reserve specific upstream transmission opportunities for specific real-time traffic flows. The CMTS MUST provide fixed size data grants at periodic intervals to the Service Flow. The CM MUST use only unsolicited data grants for upstream transmission. The key service information elements are the Unsolicited Grant Size, the Nominal Grant interval, the Tolerated Grant Jitter and the Request/Transmission Policy. (Refer to Appendix M)

The Unsolicited Grant Synchronization Header (UGSH) in the Service Flow EH Element (refer to Section 6.2.6.3.2) is used to pass status information from the CM to the CMTS regarding the state of the UGS Service Flow. The most significant bit of the UGSH is the Queue Indicator (QI) bit. The CM MUST set this flag once it detects that this Service Flow has exceeded its transmit queue depth. Once the CM detects that the Service Flow's transmit queue is back within limits, it MUST clear the QI flag. The flag allows the CMTS to provide for long term compensation for conditions such as lost maps or clock rate mismatch's by issuing additional grants.

The CMTS MUST NOT grant in excess of active Grants per Interval, excluding the case when the QI bit in the UGSH is set. In this case, the CMTS MAY grant up to 1% additional bandwidth for clock rate mismatch compensation. The CMTS policing of the Service Flow remains unchanged.

8.2.2 Real-Time Polling Service

The intent of rtPS is to reserve upstream transmission opportunities for real-time traffic flows (such as voice over IP). These Service Flows receive periodic transmission opportunities regardless of network congestion, but these Service Flows release their transmission reservations to other Service Flows when inactive. The CMTS polls rtPS SIDs on a periodic basis (typically on the order of tens of milliseconds or less) for current upstream traffic usage through the use of unicast request opportunities.

The CMTS MUST provide periodic unicast request opportunities. The CM MUST use only unicast request opportunities in order to obtain upstream transmission opportunities (the CM MUST use unsolicited data grants for upstream transmission as well). The key service information elements are the Nominal Polling Interval, the Tolerated Poll Jitter and the Request/Transmission Policy.

8.2.3 Unsolicited Grant Service with Activity Detection

The intent of UGS-AD is to reserve upstream transmission opportunities for real-time traffic flows (such as voice-over-IP with silence suppression) in a way that takes advantage of the reduced latency of Unsolicited Grants, when active, but has the efficiency of Real-Time Polling when inactive.

The CMTS MUST provide periodic unicast grants, when the flow is active, but MUST revert to providing periodic unicast request opportunities when the flow is inactive. [The CMTS can detect flow inactivity by detecting unused grants. However, the algorithm for detecting a flow changing from an active to an inactive state is dependent on the CMTS implementation] The CM MUST use only unicast request opportunities in order to obtain upstream transmission opportunities. The CM MUST use unsolicited data grants for upstream transmission as well. The key service parameters are the Nominal Polling Interval, the Tolerated Poll Jitter, the Nominal Grant Interval, the Tolerated Grant Jitter, the Unsolicited Grant Size and the Request/Transmission Policy.

In UGS-AD service, when restarting UGS after an interval of RTP, the CMTS SHOULD provide additional grants in the first (and/or second) grant interval such that the CM receives a total of one grant for each grant interval from the time the CM requested restart of UGS, plus one additional grant. (Refer to Appendix M) Because the Service Flow is provisioned as a UGS flow with a specific grant interval and grant size, when restarting UGS, the CM MUST NOT request a different sized grant than the already provisioned UGS flow. As with any Service Flow, changes can only be requested with a DSC command.¹

The Service Flow Extended Header Element allows for the CM to dynamically state how many grants per interval are required to support the number of flows with activity present. In UGS/AD, the CM MAY use the Queue Indicator Bit in the UGSH. The remaining seven bits of the UGSH define the Active Grants field. This field defines the number of grants within a Nominal Grant Interval that this Service Flow currently requires. The CM MAY indicate from 0 to Grants Per Interval grants active per Nominal Grant Interval. This field allows the CM to signal to the CMTS to dynamically adjust the number of grants per interval that this UGS Service Flow is actually using. The CM MUST NOT request more than the number of Grants per Interval in the ActiveQoSParameterSet.

8.2.4 Non-Real-Time Polling Service

The intent of nrtPS is to set aside upstream transmission opportunities for non-real-time traffic flows (such as high bandwidth FTP). These flows receive some transmission opportunities during network congestion. The CMTS typically polls nrtPS SIDs on an (periodic or non-periodic) interval on the order of one second or less.

The CMTS MUST provide timely unicast request opportunities. The CM MUST use both contention request and unicast request opportunities in order to obtain upstream transmission opportunities (the CM MUST use unsolicited data grants for upstream transmission as well). The key service elements are Nominal Polling Interval, Reserved Minimum Traffic Rate, Maximum Sustained Traffic Rate, Request/Transmission Policy and Traffic Priority.²

8.2.5 Best Effort Service

The intent of the Best Effort (BE) service is to provide efficient service to best effort traffic. With BE service, the CM MUST use all (contention and unicast) request opportunities³, and all unicast data transmission opportunities. The CM MAY use contention data opportunities as appropriate. The key service elements are the Minimum Reserved Traffic Rate, the Maximum Sustained Traffic Rate, and the Traffic Priority.⁴

8.2.6 Other Services

8.2.6.1 Committed Information Rate (CIR)

A Committed Information Rate (CIR) service can be defined a number of different ways. For example, it could be configured by using a Best Effort service with a Reserved Minimum Traffic Rate or a nrtPS with a Reserved Minimum Traffic Rate.

1. Sentenc added 06/22/99 per rfi-n-99043 ew

2. Edited 06/22/99 per rfi-n-99043 ew

3. With appropriate deference for contention request opportunities. Refer to 7.4.1.

4. Edited 06/22/99 per rfi-n-99043 ew

8.3 Fragmentation

Fragmentation is an upstream CM “modem capability”. The CMTS MUST enable or disable this capability on a per-modem basis with a TLV in the Registration Response. The per-modem basis provides compatibility with DOCSIS 1.0 CMs. Once fragmentation is enabled for a DOCSIS 1.1 modem, fragmentation is enabled on a per-Service Flow basis via the Request/Transmission Policy Configuration Settings. When enabled for a Service Flow, fragmentation is initiated by the CMTS when it grants bandwidth to a particular CM with a grant size that is smaller than the corresponding bandwidth request from the CM. This is known as a **Partial Grant**.

8.3.1 CM Fragmentation Support

Fragmentation is essentially encapsulation of a portion of a MAC Frame within a fixed size fragmentation header and a fragment CRC. Concatenated PDUs, as well as single PDUs, are encapsulated in the same manner. Baseline Privacy, if enabled, is performed on each fragment as opposed to the complete original MAC frame.

The CM MUST perform fragmentation according to the flow diagram in Figure 8-8. The phrase “untransmitted portion of packet” in the flow diagram refers to the entire MAC frame when fragmentation has not been initiated and to the remaining untransmitted portion of the original MAC frame when fragmentation has been initiated.

8.3.1.1 Fragmentation Rules:

1. Any time fragmentation is enabled and the grant size is smaller than the request, the CM MUST fill the partial grant it receives with the maximum amount of data (fragment payload) possible accounting for fragmentation overhead and physical layer overhead.
2. The CM MUST send up a piggyback request any time there is no later grant or grant pending for that SID in MAPs that have been received at the CM.
3. If the CM is fragmenting a frame¹, any piggyback request MUST be made in the BPI EHDR portion of the fragment header.
4. In calculating bandwidth requests for the remainder of the frame (concatenated frame, if concatenated) that has been fragmented, the CM MUST request enough bandwidth to transmit the entire remainder of the frame plus the 16-byte fragment overhead and all associated physical layer overhead.
5. If the CM does not receive a grant or grant pending within the ACK time of sending a request, the CM MUST backoff and re-request for the untransmitted portion of the frame until the bandwidth is granted or the CM exceeds its retry threshold.
6. If the CM exceeds its retry threshold while requesting bandwidth, the CM discards whatever portion of the frame was not previously transmitted.
7. The CM MUST set the F bit and clear the L bit in the first fragment of a frame.
8. The CM MUST clear the F and L bits in the fragment header for any fragments that occur between the first and last fragments of a frame.
9. The CM MUST set the L bit and clear the F bit in the last fragment of a frame.
10. The CM MUST increment the fragment sequence number sequentially for each fragment of a frame transmitted.
11. If a frame is to be encrypted and the frame is fragmented, the frame is encrypted only at the fragment layer with encryption beginning immediately after the fragment header HCS and continuing through the fragment CRC.
12. Frames sent in immediate data (request/data) regions MUST NOT be fragmented.

1. Note, ‘frame’ always refers to either frames with a single Packet PDU or concatenated frames.

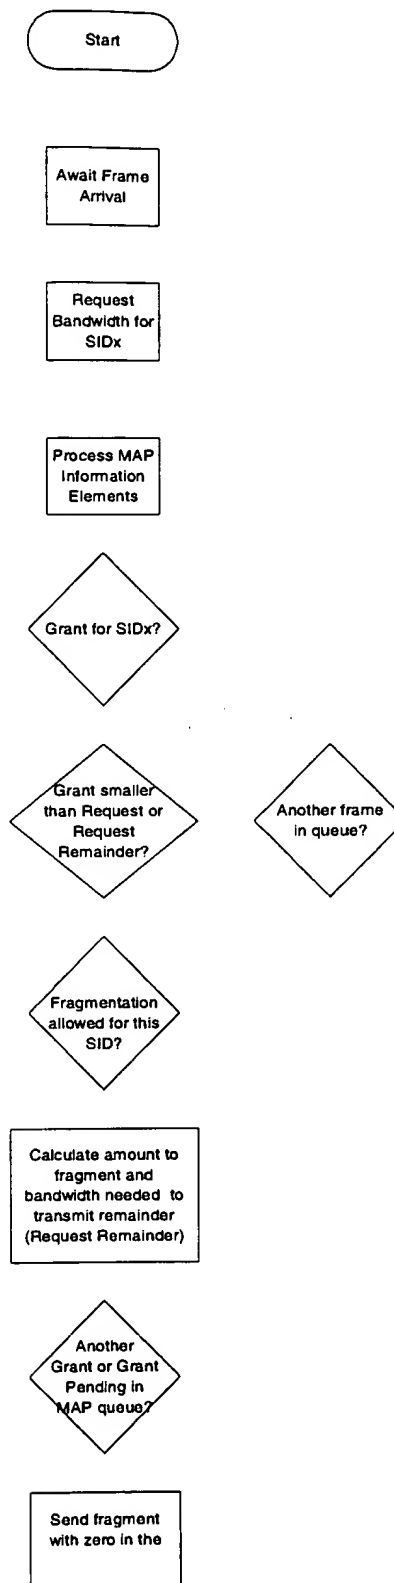


Figure 8-8. CM Fragmentation Flowchart

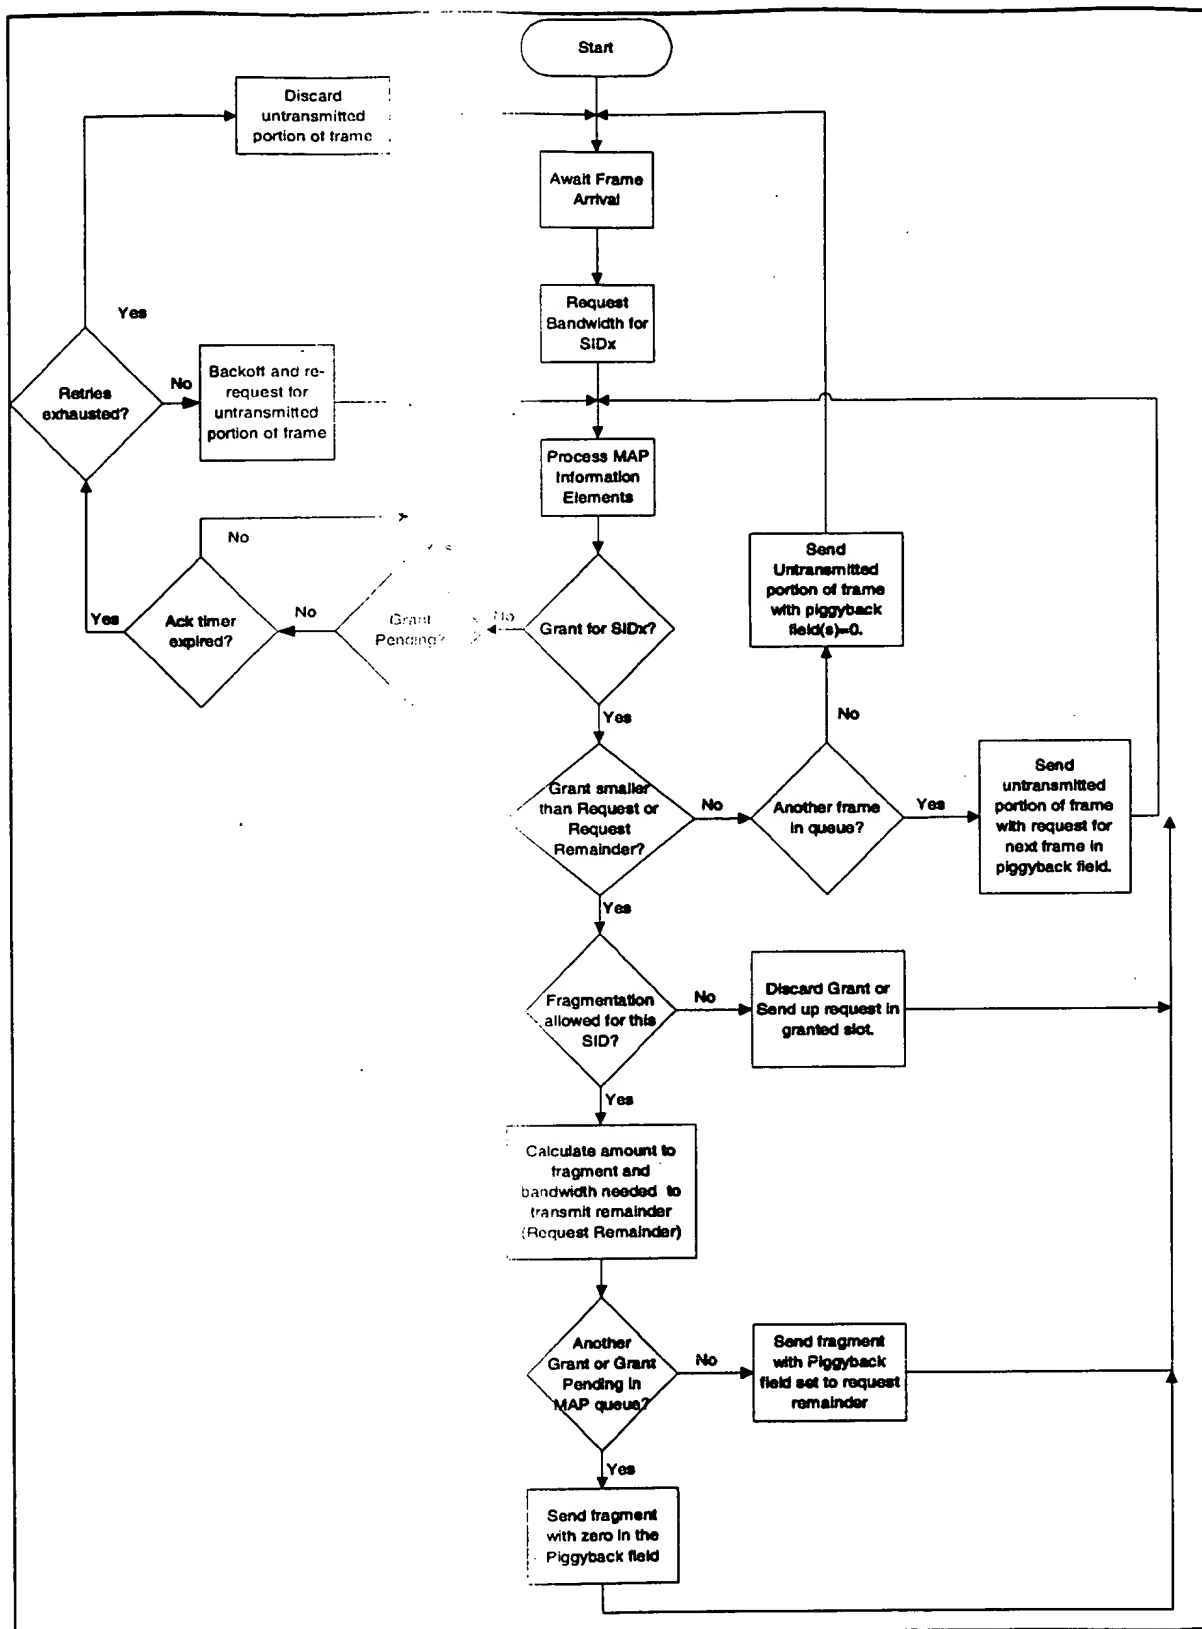


Fig. 4-6-8. CM Fragmentation Flowchart

8.3.2 CMTS Fragmentation Support

At the CMTS, the fragment is processed similarly to an ordinary packet with the exception that the baseline privacy encryption starts just after the fragment header as opposed to being offset by 12 bytes.

The CMTS has two modes it can use to perform fragmentation. The Multiple Grant Mode assumes that the CMTS retains the state of the fragmentation. This mode allows the CMTS to have multiple partial grants outstanding for any given SID. The Piggybacking Mode assumes the CMTS does NOT retain any fragmentation state. Only one partial grant is outstanding, so that the CM inserts the remaining amount into the Piggyback field of the fragment header. The type of mode being used is determined by the CMTS. In all cases, the CM operates with a consistent set of rules.

8.3.2.1 Multiple Grant Mode

A CMTS MAY support Multiple Grant Mode for performing fragmentation.

Multiple Grant Mode allows the CMTS to break a request up into two or more grants in a single or over successive maps and it calculates the additional bandwidth required in the remaining partial grants to satisfy the request. In Multiple Grant Mode, if the CMTS cannot grant the remainder in the current MAP, it MUST send a grant pending (zero length grant) in the current MAP and all subsequent MAPs to the CM until it can grant additional bandwidth. If there is no grant pending in subsequent maps, the CM MUST re-request for the remainder. This re-request mechanism is the same as that used when a normal REQ does not receive a grant or grant pending within the ACK time.

If a CM receives a grant pending IE along with a fragment grant, it MUST NOT piggyback a request in the extended header of the fragment transmitted with the grant.

In the case where the CM misses a grant and requests the remaining bandwidth, the CMTS MUST recover without dropping the frame.

Due to the imprecision of the mini-slot to byte conversion process, the CMTS MUST make sure that any fragment payload remainder is greater than one slot (i.e. the imprecision amount). Failure to do this may cause the CMTS to issue a grant that is not needed after the CM has completed transmission of the fragment payload remainder using a previous partial grant. This may cause the CM to get out of sync with the CMTS by inadvertently starting a new fragmentation.

8.3.2.2 Piggyback Mode

A CMTS MAY support Piggyback Mode for performing fragmentation.

If the CMTS does not put another partial grant pending in the MAP in which it initiates fragmentation on a SID, the CM MUST automatically request for the remainder. The CM calculates how much of a frame can be sent in the granted bandwidth and requests the remainder to send it. The CM utilizes the piggyback field in the fragment extended header to request the remainder necessary to transfer the remainder of the frame. Since the CMTS did not indicate a multiple grant in the first fragment MAP, the CM MUST keep track of the remainder to send. The request length, including physical and fragmentation overhead, for the remainder of the original frame is inserted into the piggyback request field in the fragmentation header.

If the fragment HCS is correct, the piggyback request, if present, is passed on to the bandwidth allocation process while the fragment itself is queued for reassembly. Once the complete MAC Frame is reassembled,

1. Sentence clarified 06/22/99 per rfi-n-99-0024

any non-privacy extended headers in the packet. If the packet HCS is correct, and the packet is forwarded to the appropriate destination.

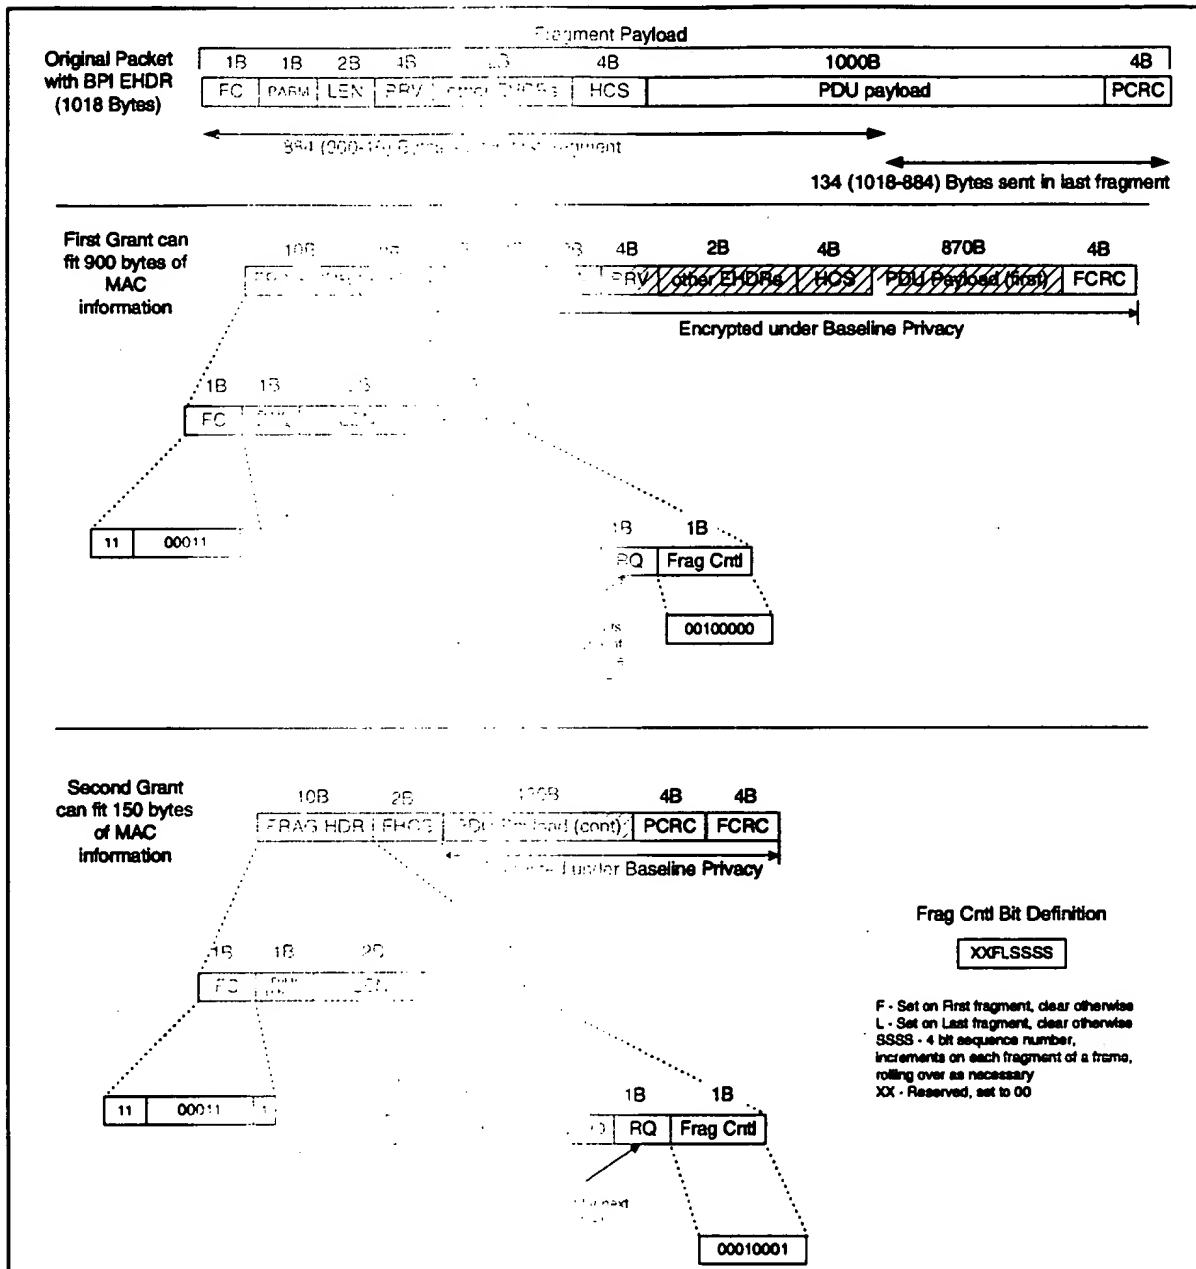
8.3.3 Fragmentation Example

8.3.3.1 Single Packet Fragmentation

Refer to Figure 8-8. Assume that fragmentation has been enabled for a given SID.

1. **(Requesting State)**- CM wants to transmit a 1018 byte packet. CM calculates how much physical layer overhead (POH) is required and requests the appropriate number of minislots. CM makes a request in a contention region. Go to step 2.
2. **(Waiting for Grant)**- CM monitors for a grant or grant pending for this SID. If the CM's ACK time expires before the CM receives a grant or grant pending, the CM retries requesting for the packet until the retry count is exhausted - then the CM times up on that packet. Go to step 3.
3. **(First Fragment)**- Prior to giving up in step 2, the CM sees a grant for this SID that is less than the requested number of minislots. The CM calculates how much MAC information can be sent in the granted number of minislots using the specified fragmentation overhead. In the example in Figure 8-9, the first grant can hold 900 bytes after subtracting the POH. Since the fragmentation overhead (FRAG HDR, FHCS, and FCRC) is 16 bytes, 884 bytes of the original packet can be sent in the fragment. The CM creates a fragment composed of the FRAG HDR, FHCS, 884 bytes of the original packet, and an FCRC. The CM marks the fragment as first and prepares to send the fragment. Go to step 4.
4. **(First Fragment, multiple grants)**- The CM looks to see if there are any other grants or grant pendings enqueued for this SID. If so, the CM sets the piggyback field in the FRAG HDR set to zero and awaits the next grant to roll around. --- go to step 6. If there are not any grants or grant pendings, go to step 5.
5. **(First Fragment, single grant)**- If there are no other grants or grant pendings for this SID in this MAP, the CM calculates the fragmentation overhead, and adds this amount to the piggyback field of the FRAG HDR. The CM then sends the fragment and starts its ACK timer for the piggyback request. In the example in Figure 8-9, the CM sends a request for enough minislots to hold the POH plus 150 bytes (1018-884+16). Go to step 6.
6. **(Waiting for Grant)**- The CM is monitoring for a grant for the next fragment. If the CM's ACK timer expires while waiting on this grant, the CM sends up a request for enough minislots to send the remainder of the fragmented packet, including the fragmentation overhead, and physical layer overhead. Go to step 7.
7. **(Receives next fragment grant)**- Following up in step 6, the CM sees another grant for this SID. The CM checks to see if the grant size is enough to hold the remainder of the fragmented packet, including the fragmentation overhead and physical layer overhead. If so, go to step 10. If not, go to step 8.
8. **(Middle Fragment, multiple grants)**- If the remainder of the packet (plus overhead) will not fit in the grant, the CM encapsulates this portion of the packet as a middle fragment. The CM then looks for any grants or grant pendings enqueued for this SID. If either are present, the CM sets the piggyback field in the FRAG HDR set to zero and awaits the next grant to roll around. --- go to step 6. If there are not any grants or grant pendings, go to step 9.
9. **(Middle Fragment, single grant)**- If there are no other grants or grant pendings for this SID in this MAP, the CM calculates how many minislots are required to send the remainder of the fragmented packet, including the fragmentation overhead and physical layer overhead, and inserts this amount into the piggyback field of the FRAG HDR. The CM then sends the fragment and starts its ACK timer. Go to step 6.

10. (Last Fragment) - The CM encapsulates the remainder of the packet as a last fragment. If there is no other packet enqueued or there is another packet of grant pending enqueued for this SID, the CM places a zero in the REQ field of the FRAG HDR. If there is another packet enqueued with no grant of grant pending, the CM calculates the number of milliseconds until the next packet and places this number in the REQ field in the FRAG HDR. The CM then transmits the packet. Go to step 11. In the example in Figure 8-9, the grant is large enough to hold the remainder of the packet.
11. (Normal operation) - The CM then enters a normal operation of waiting for grants and requesting for packets. If at any time fragmentation occurs and a grant arrives that is smaller than the request, the fragmentation process starts again.



8.3.3.2 Concatenated Packet Fragmentation

After the CM creates the concatenated packet, the CM treats the concatenated packet as a single PDU. Figure 8-10 shows an example of a concatenated packet broken into 3 fragments. Note that the packet is fragmented without regard to the packet boundaries of the concatenated packet.




Figure 8-10 Concatenated Packet Example

8.4 Payload Header Suppression

The overview section explains the primary use of the lead Header Suppression. The subsequent sections explain the signaling for initialization, opened, and closed. Finally, specific upstream and downstream examples are given. The following details:

		for Suppression Definitions
PHS	Payload Header	Suppressing an initial byte string at the sender and restoring the byte string at the receiver.
PHS Rule	Payload Header	A set of TLV's that apply to a specific PHS Index.
PHSF	Payload Header Field	The value of the suppressed byte string.
PHSI	Payload Header Index	An 8-bit value which references the suppressed byte string.
PHSM	Payload Header Mask	A bit mask which indicates which bytes in the PHSF to suppress, and which bytes to not suppress.
PHSS	Payload Header Size	The length of the Suppressed Field in bytes.
PHSV	Payload Header Verify	A flag which tells the sending entity to verify all bytes which are to be suppressed.

8.4.1 Overview

In Payload Header Suppression, the transmission of one or more of the payload headers following the Extended Header field is suppressed by the sending entity and recovered by the receiving entity. In the upstream, the sending entity is the CM and the receiving entity is the CMTS. In the downstream, the sending entity is the CMTS and the receiving entity is the CM. The CM or CMTS includes a Payload Header Suppression Index (PHSI) which references the Payload Header Suppression Table.

Although PHS may be used with the Unsolicited Grant Service (UGS) Service Class, it was designed for use with UGS because UGS packets are always the same length. With UGS, PHS will always produce a fixed length compressed packet header.¹

The sending entity uses the PHSF to map packets to a Service Flow. The Classifier uniquely maps packets to its associated Payload Header Number (PHN). The receiving entity uses the Service Identifier (SID)² and the PHSI to restore the PHSF. Once a PHSF has been assigned to a PHSI, it cannot be changed. To change the value of a PHSF on a Service Flow, a new Payload Header Number and a new Classification Rule must be defined, the old rule is removed from the Service Flow, and the old Classifier is deleted. If the old Classifier is deleted, any associated PHS rule MUST also be deleted.³

PHS has a PHSV option to suppress the following bytes before suppressing it. PHS also has a PHSM option to allow select bytes to be suppressed. The following bytes are not suppressed: IP sequence numbers, and still others.

1. Paragraph rephrased to read:

2. **No SID is necessary** since it will be sufficient since it applies to all downstream Service Flows on a C

3. Sentence ad'

PHS rules are consistent with the FCC's rules. Requests and grants of bandwidth are specified after suppression has been completed. For **Unsolicited Grant Services**, the grant size is chosen with the Unsolicited Grant Size TLM. The packet size of the request may be equal to or less than the grant size.

The CMTS MUST assign all PHS values. It assigns all SID values. Either the sending or the receiving entity MAY send the PHS in the PHS field. The extension allows for pre-configured headers, or for higher level signaling protocols, or for a separate signaling channel to establish cache entries. PHS is intended for unicast service, and is not intended for multicast service.

It is the responsibility of the higher-layer service entity to generate a PHS Rule which uniquely identifies the higher-layer service entity to the lower-layer service entity. It is the responsibility of the higher-layer service entity to maintain the PHS Rule constant from packet to packet for the duration of the Active Service Flow.

8.4.2 Example 8.4.2

- A Classifier can be configured to identify a media flow uniquely defines a Voice-over-IP (VoIP) flow by specifying the Source Port, UDP Destination Port, the Service Flow Reference, and a PHSI value. This Classifier provides a PHSI value which identifies the media flow. The first 14 bytes of payload header are verified and suppressed, and a 2 byte extension is added to every packet in that media flow.
- A Classifier can be configured to identify a media flow, of which 90% match the PHSR. Verification is enabled for the media flow where every so often compression resets are done. The compression algorithm would allow variable bandwidth, and only 90% of the packets would be compressed. Since the existence of the PHSI extended header will be verified, the backup at the receiving entity will always yield the correct result.
- A Classifier can be configured to identify all IP packets by specifying the Service Flow Reference, the Ethernet Type, and the Ethernet Header. This Classifier provides no verification by the sending entity. In this example, the Ethernet Header is verified, but it will not require the first 14 bytes of the Ethernet header. The Source Address or Destination Address may vary. The CM removes the Ethernet Header (without verifying their contents) and forwards the frame.

8.4.3 C

To clarify or modify the protocol, the following procedure shall be followed. The following procedure shall be followed for one potential implementation. CM and CMTS implementers shall not be required to implement the protocol for Suppression in any manner as long as the protocol specified in this document is followed. The following procedure shall be followed.

The CM applies its list of Classifier rules. A match of the PHS Rule. The PHS Rule provides PHSF, PHSI, PHSM, PHSS, and PHSE. The PHSF and PHSI will verify the Upstream Suppression Field in the packet. The PHSM will verify the bytes in the Upstream Suppression Field except the PHSI into the PHS_Parm field of the Service Flow EH. The PHSS will verify the Flow.

When the packet is received, the CMTS will determine the associated SID either by internal means or from the PHSI. The CMTS uses the SID and the CMTS Extended Header. The CMTS uses the SID and the PHSI to locate the packet in the cache. The CMTS then reassembles the packet and then proceeds with normal packet processing. The CMTS then removes the first 4 bytes from the PHSF. If verification was enabled, then

the PHSF bytes will equal the original PHSF bytes. If verification was not enabled, then there is no guarantee that the PHSF bytes will match the original PHSF bytes.

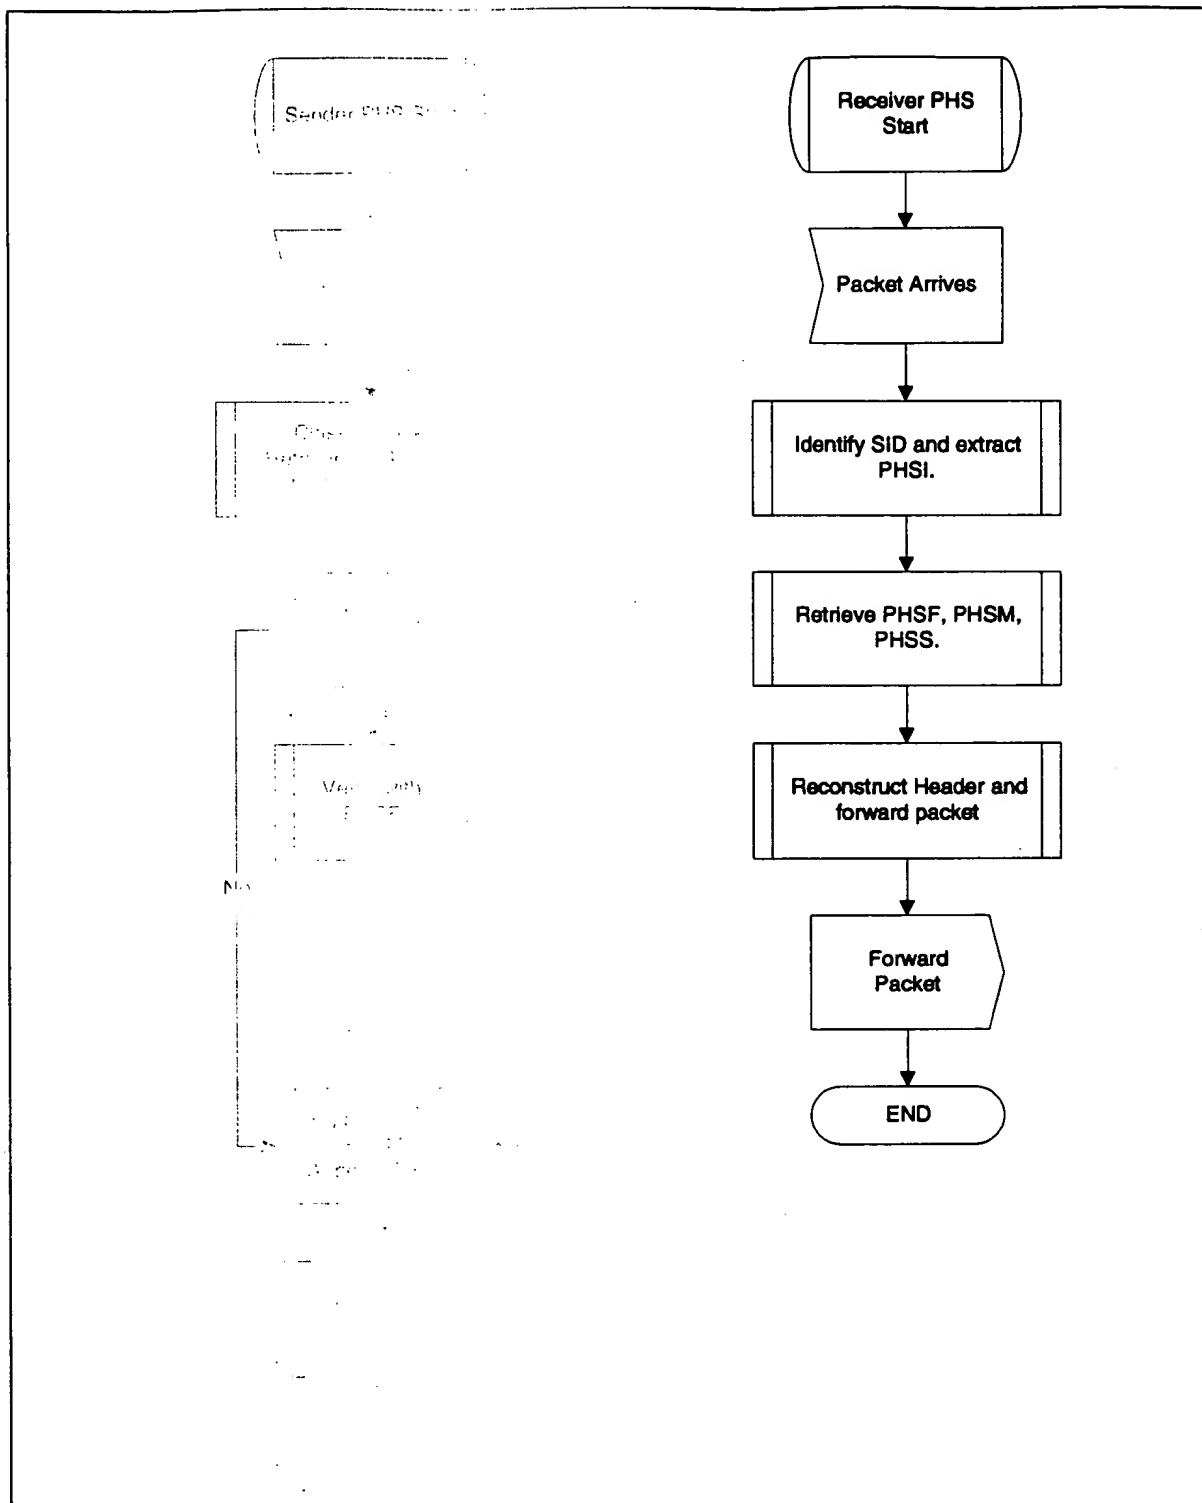



Figure 2-11: Inband Header Suppression Operation

A similar operation occurs for the Downstream. The CMTS applies its list of Classifiers. A match of the Classifier will result in a Downstream Classifier being associated with a PHS Rule. The PHS Rule provides PHSF, PHSI, PHSM, PHSS, and PHSV. If PHSV is set on a PHS Rule, the CMTS will verify the Downstream Suppression Field in the packet with the PHSF. If they match, the CMTS will suppress all the bytes in the Downstream Suppression Field except the bytes masked by PHSM. The CMTS will then insert the PHSI into the PHS_Parm field of the Service Flow EH Element, and queue the packet into the Downstream Service Flow.

The CM will receive the packet based upon the Ethernet Destination Address filtering. The CM then uses the PHSI to look up PHNR, PHSM, and PHSV. The CM reassembles the packet and then proceeds with normal packet processing.

Figure 8-12 demonstrates the packet suppression and restoration when using PHS masking. Masking allows only bytes which are in the PHSM to be suppressed. The PHSF and PHSS span the entire Suppression Field, included segments of the packet.



8.4.4 Signaling

Payload Header Signaling

The following objects are used in the signaling process:

- Service Flow
- Classifier
- PHS Rule

These three objects are created by the CMTS and may be created simultaneously.

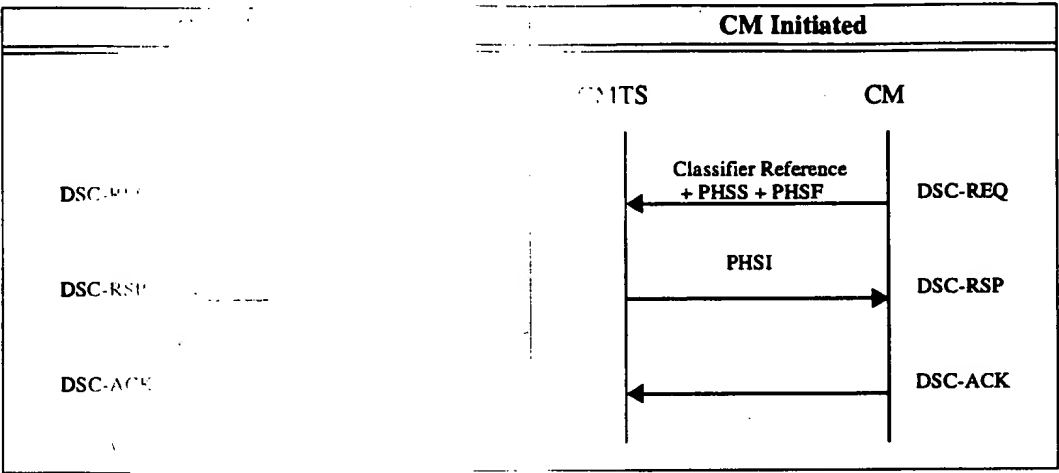
PHS Rules are used to define the suppression and restoration of messages. The CMTS MUST define the PHSI when the PHS Rule is used for suppression or restoration. The CM or CMTS MAY define the PHSS and PHSM.

Figure 8-13 shows the structure of a PHS Rule.¹

It is possible to partially specify a rule when a Service Flow is created. As an example, if the PHSF is suppressed will be known, the PHSS will be known and would be provided in the "Set PHS Rule" DSC-REQ. The PHSS is provided in the registration request or a DSC-REQ. The PHSS uniquely identify the PHS rule. The PHSS

particular the size of the rule) at the time a Service Flow is provisioned the header fields to be suppressed. The PHSS (e.g., IP addresses, UDP port numbers, etc.) may not be known at the time the Service Flow is provisioned. The PHSS is provided in the "Set PHS Rule" DSC-REQ as part of the activation of the Service Flow (using the PHSS defined in more than one step, each step, whether it is a registration request or a DSC-REQ). The PHSS is provided in the Service Flow ID (or reference) and a PHS index to

Figure 8-10. CM Initiated Header Suppression Signaling Example



1. Sentence fragment

9913

1. paragraph added

8.4.5 Payload Header Suppression Examples

8.4.5.1 Upstream Service Flow

A Service Class with the Service Flow ID of "G711-US-UGS-HS-42" is established which is intended for G711 VoIP traffic in the upstream channel. When Classifiers are added to the flow, a PHSS value of 42 is included in the MAC Header Checksum. The MAC Header states that the first 42 bytes following the MAC Extended Header will be suppressed, and restored. In this example, the Service Class is configured such that the first 42 bytes of the packet will not be verified, and will be discarded since it is not within the Grant Size. (Refer to C.2.2.6.3)

Figure 8-14 shows the structure of the upstream with and without Payload Header Suppression. An RTP Voice packet is used as a specific example to demonstrate efficiency.

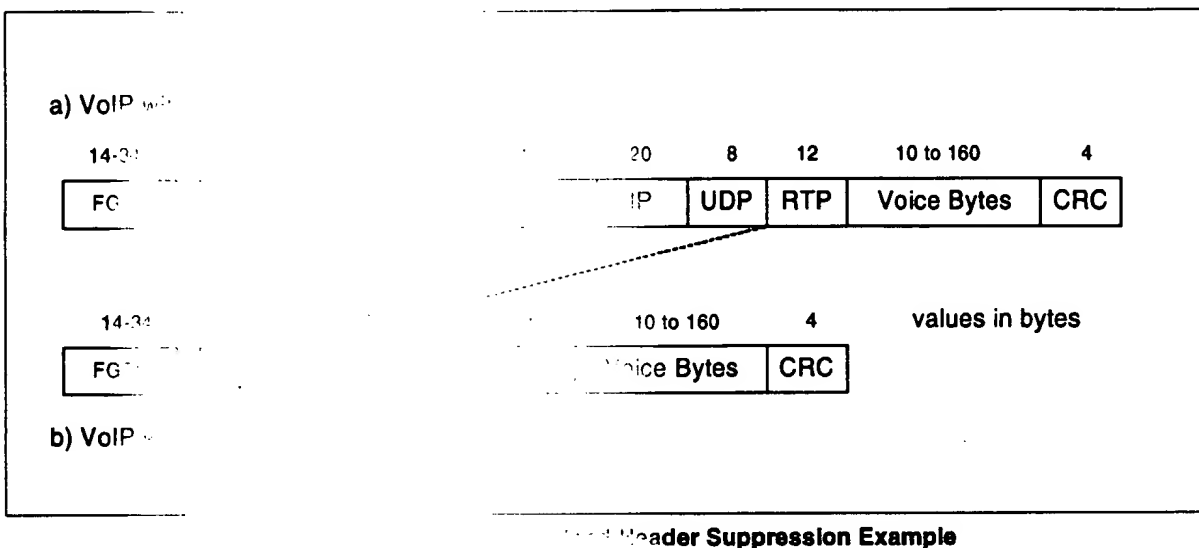


Figure 8-14 illustrates the structure of the upstream channel. The beginning of the frame represents the physical layer preamble, and stuffing bytes. Stuffing bytes occur in the last code word of the frame. Next is the MAC layer overhead including the 6 byte Ethernet Header, and the 4 byte Ethernet CRC. The UDP header, and a 12 byte RTP header. The voice payload is the compression algorithm used.

Figure 8-14 illustrates the structure of the upstream channel with Payload Header Suppression enabled. In the upstream, Payload Header Suppression is enabled. The 14 byte Ethernet header, the 6 byte MAC Header Checksum. The 14 byte Ethernet header, the 6 byte MAC Header Checksum, and the 2 byte PHS Extended Header element are suppressed, and a 2 byte PHS Extended Header element is included. In the example of an established VoIP connection, these fields are redundant.

8.4.5.2 Downstream Service Flow

A Service Class with the Service Flow ID of "G711-US-UGS-HS-30" is established which is intended for G711 VoIP traffic in the downstream channel. When Classifiers are added to the flow, a PHSS value of 30 is included in the MAC Header Checksum. The MAC Header states that the first 30 bytes following the MAC Extended Header will be suppressed, and restored. In this example, the Service Class is configured such that the first 30 bytes of the packet must be verified, suppressed and restored. The packet will not have its header suppressed but will be transmitted in the downstream channel.

Figure 8-15 shows the encapsulation for downstream with and without Payload Header Suppression. An RTP Voice over IP Payload is used as a specific example to demonstrate efficiency.




Figure 8-15a shows the encapsulation for downstream with Payload Header Suppression enabled. The Layer 2 overhead includes the 6 byte MAC header (2 byte Destination Address, 2 byte Source Address, and 2 byte MAC Header Checksum). The Layer 3 VoIP payload uses a 2 byte UDP header, a 12 byte RTP header, and a variable length voice payload. The voice payload is variable and depends on the codec used.

Figure 8-15b shows the encapsulation for downstream with Payload Header Suppression disabled. In the downstream, Payload Header Suppression is not used. The MAC Header Checksum is retained, so that the CM may filter and receive the packet. The Destination Address and the 8 byte UDP header have been suppressed, and the 2 byte RTP header has been added, for a net reduction of 28 bytes. In this example, the 2 byte RTP header is constant from packet to packet, and are thus redundant.

9 Cable Modem

This section covers the interaction between a CM and a CMTS. The interaction can be broken down into the following categories:

9.1 CMTS

The mechanism for the CMTS to communicate with the CM (DOCSIS 1.1) is described in the following sections:

- The CMTS to the CM (non-voice)
- If valid per the Spectrum Management SYNC
- The CM to the CMTS

9.2 Cable Modem

The process for the overall flow of the cable modem (including the following sections):

The process for the cable modem (including the following sections):

- Scan
- Obtain
- Perform
- Establish
- Establish
- Trans
- Perform
- Baseline

Each CM

- A unique identifier
- Security (security)

The SDI (refer to

interaction between a CM and a CMTS. The interaction can be broken down into the following categories:

and terminal, file download, SNMP, etc.) is described in the following sections:

- In a stand-alone mode using configuration data retained in the CM
- Available storage or via another mechanism such as the Spectrum Management SYNC
- The CMTS not generate any downstream messages (including the following sections):

The process for the overall flow of the cable modem (including the following sections):

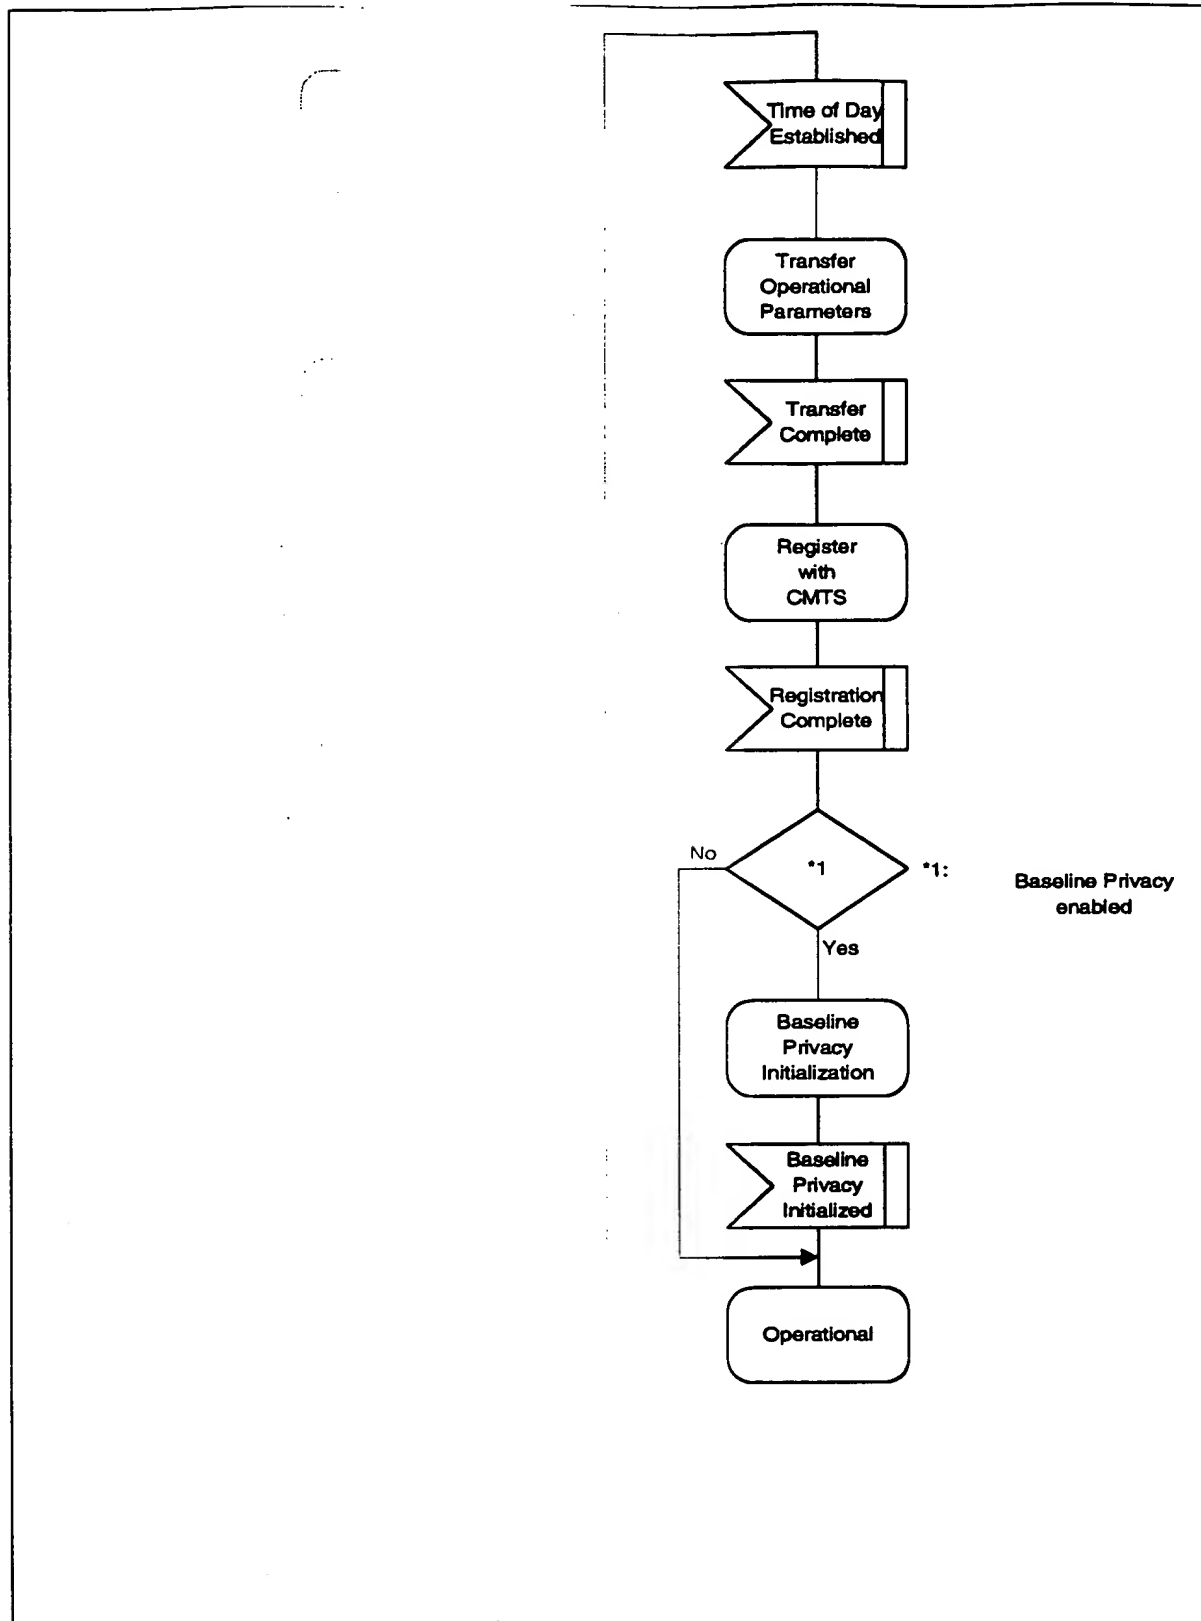
interaction with the CMTS.

(Baseline Privacy)

from the manufacturer:

- Used during the manufacturing process. This is used to identify the CM during initialization.
- (509 certificate) used to authenticate the CM to the security and provisioning servers.

The SDI used in the following figures is shown in Figure 9-2



Overview




Figure 9-2. SDL Notation

9.2.1 Scanning and Synchronization to Downstream

On initialization or after signal loss, the cable modem **MUST** acquire a downstream channel. The CM **MUST** have non-volatile storage in which the last operational parameters are stored and **MUST** first try to re-acquire this downstream channel. If this fails, it **MUST** begin to continuously scan the 6-MHz channels of the downstream frequency band of operation until it finds a valid downstream signal.

A downstream signal is considered to be valid when the modem has achieved the following steps:

- synchronization of the QAM symbol timing
- synchronization of the FEC framing
- synchronization of the MPEG packetization
- recognition of SYNC downstream MAC messages

While scanning, it is desirable to give an indication to the user that the CM is doing so.

9.2.2 Obtain Upstream Parameters

Refer to Figure 9-3. After synchronization, the CM MUST wait for an upstream channel descriptor message (UCD) from the CMTS in order to retrieve a set of transmission parameters for a possible upstream channel. These messages are transmitted periodically from the CMTS for all available upstream channels and are addressed to the MAC broadcast address. The CM MUST determine whether it can use the upstream channel from the channel description parameters.

The CM MUST collect all UCDs which are different in their channel ID field to build a set of usable channel IDs. If no channel can be found after a suitable timeout period, then the CM MUST continue scanning to find another downstream channel.

The CM MUST determine whether it can use the upstream channel from the channel description parameters. If the channel is not suitable, then the CM MUST try the next channel ID until it finds a usable channel. If the channel is suitable, the CM MUST extract the parameters for this upstream from the UCD. It then MUST wait for the next SYNC message¹ and extract the upstream mini-slot timestamp from this message. The CM then MUST wait for a bandwidth allocation map for the selected channel. It MAY begin transmitting upstream in accordance with the MAC operation and the bandwidth allocation mechanism.

The CM MUST perform initial ranging at least once per Figure 9-6. If initial ranging is not successful, then the next channel ID is selected, and the procedure restarted from UCD extraction. When there are no more channel IDs to try, then the CM MUST continue scanning to find another downstream channel.

1. Alternatively, since the SYNC message applies to all upstream channels, the CM may have already acquired a time reference from previous SYNC messages. If so, it need not wait for a new SYNC.

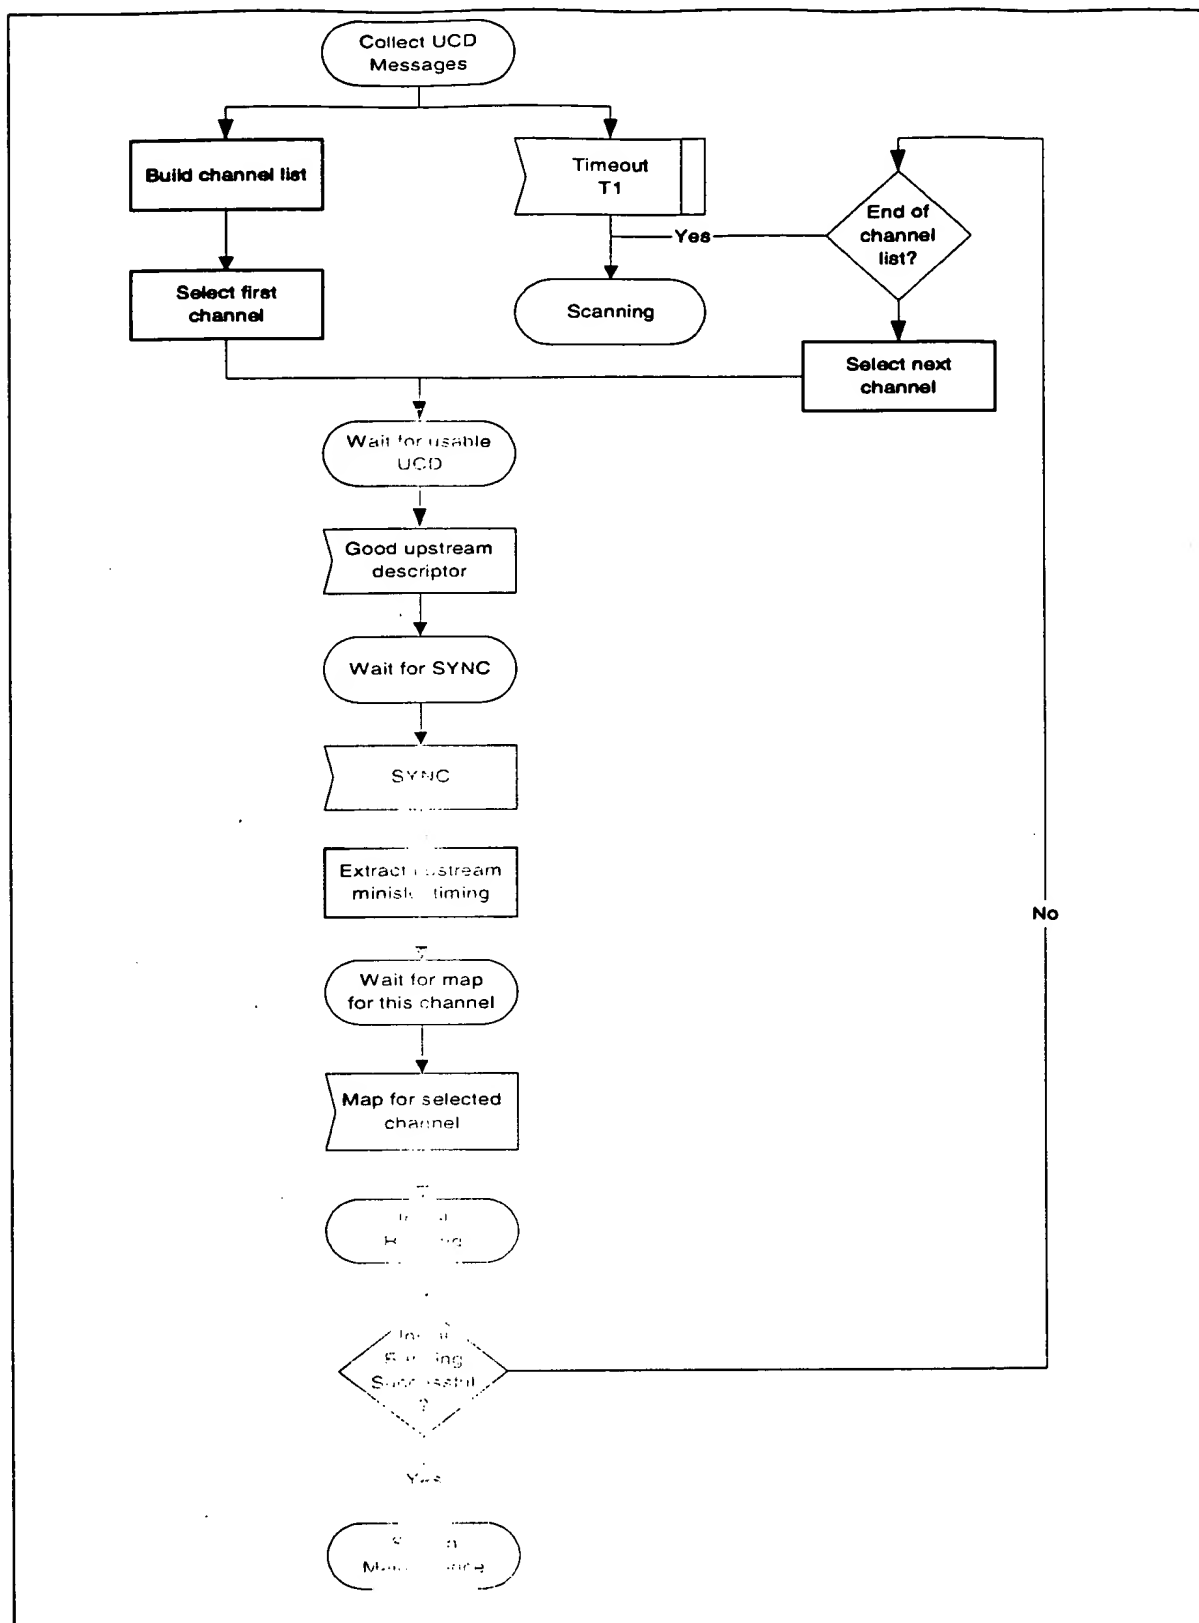


Figure 9-3. Obtaining Upstream Parameters

9.2.3 Message Flows During Scanning and Upstream Parameter Acquisition

The CMTS MUST generate SYNC and UCD messages on the downstream at periodic intervals within the ranges defined in Appendix B. These messages are addressed to all CMs. Refer to Figure 9-4.

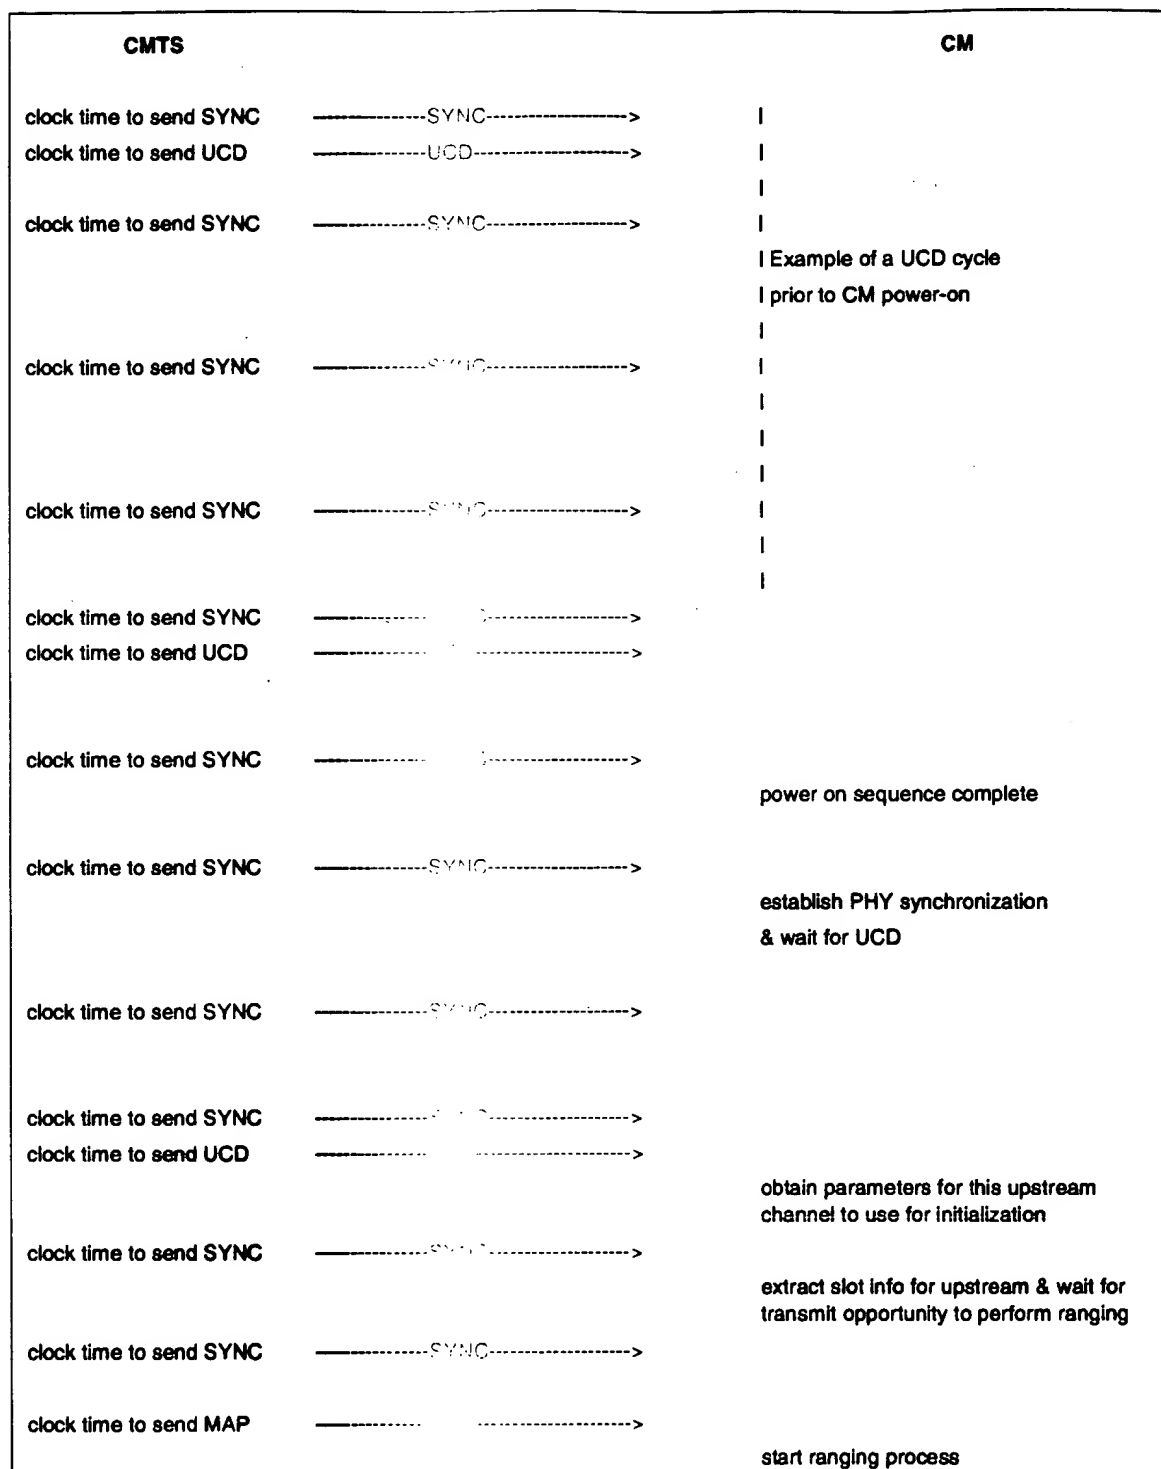


Figure 9-4. Message Flow for Downstream and Upstream Parameter Acquisition

9.2.4 Ranging and Automatic Adjustments

The ranging and adjustment process is fully defined in Section 6 and in the following sections. The message sequence chart and the finite state machines on the following pages define the ranging and adjustment process which **MUST** be followed by compliant CPEs and CMTSS. Refer to Figure 9-5 through Figure 9-8.

Note: MAPs are transmitted as described in Section 6.

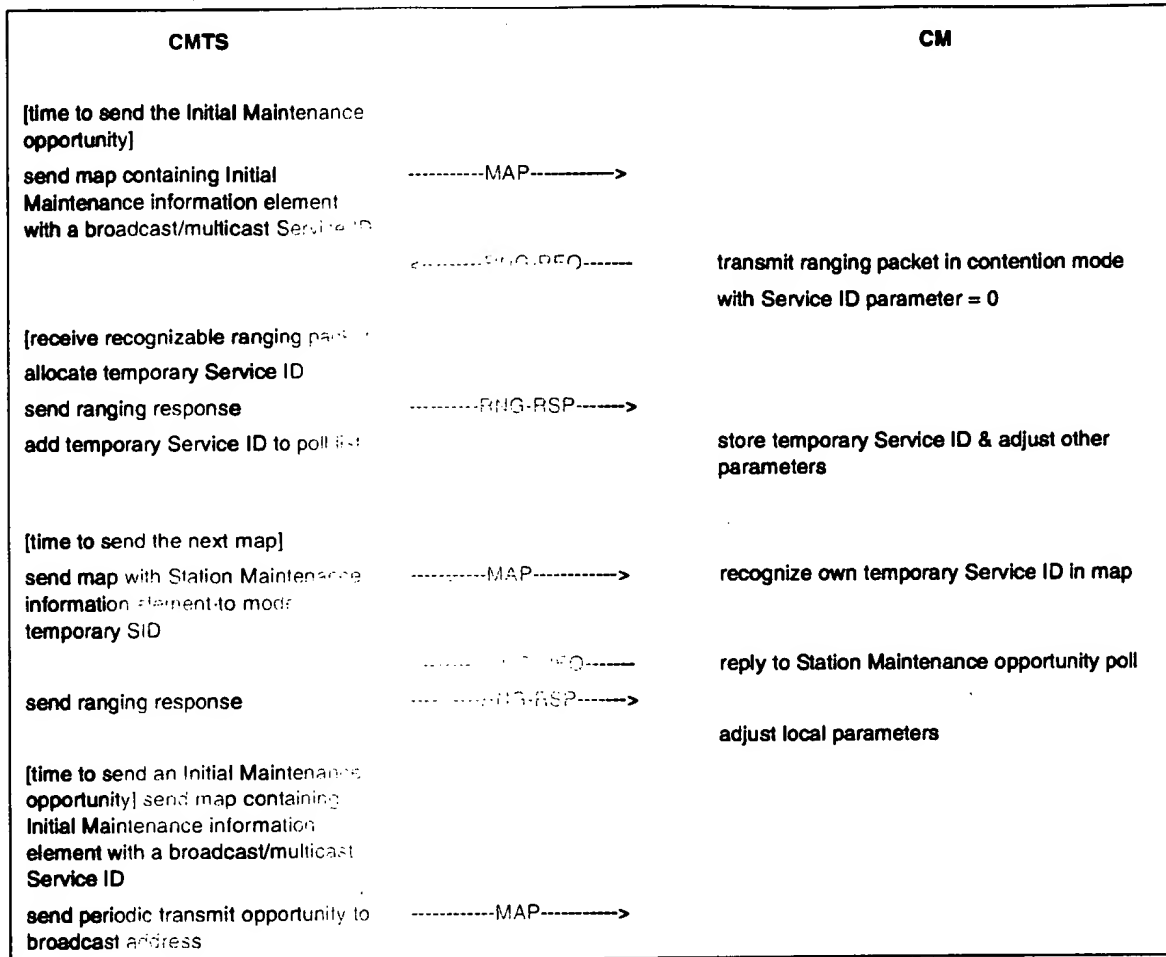
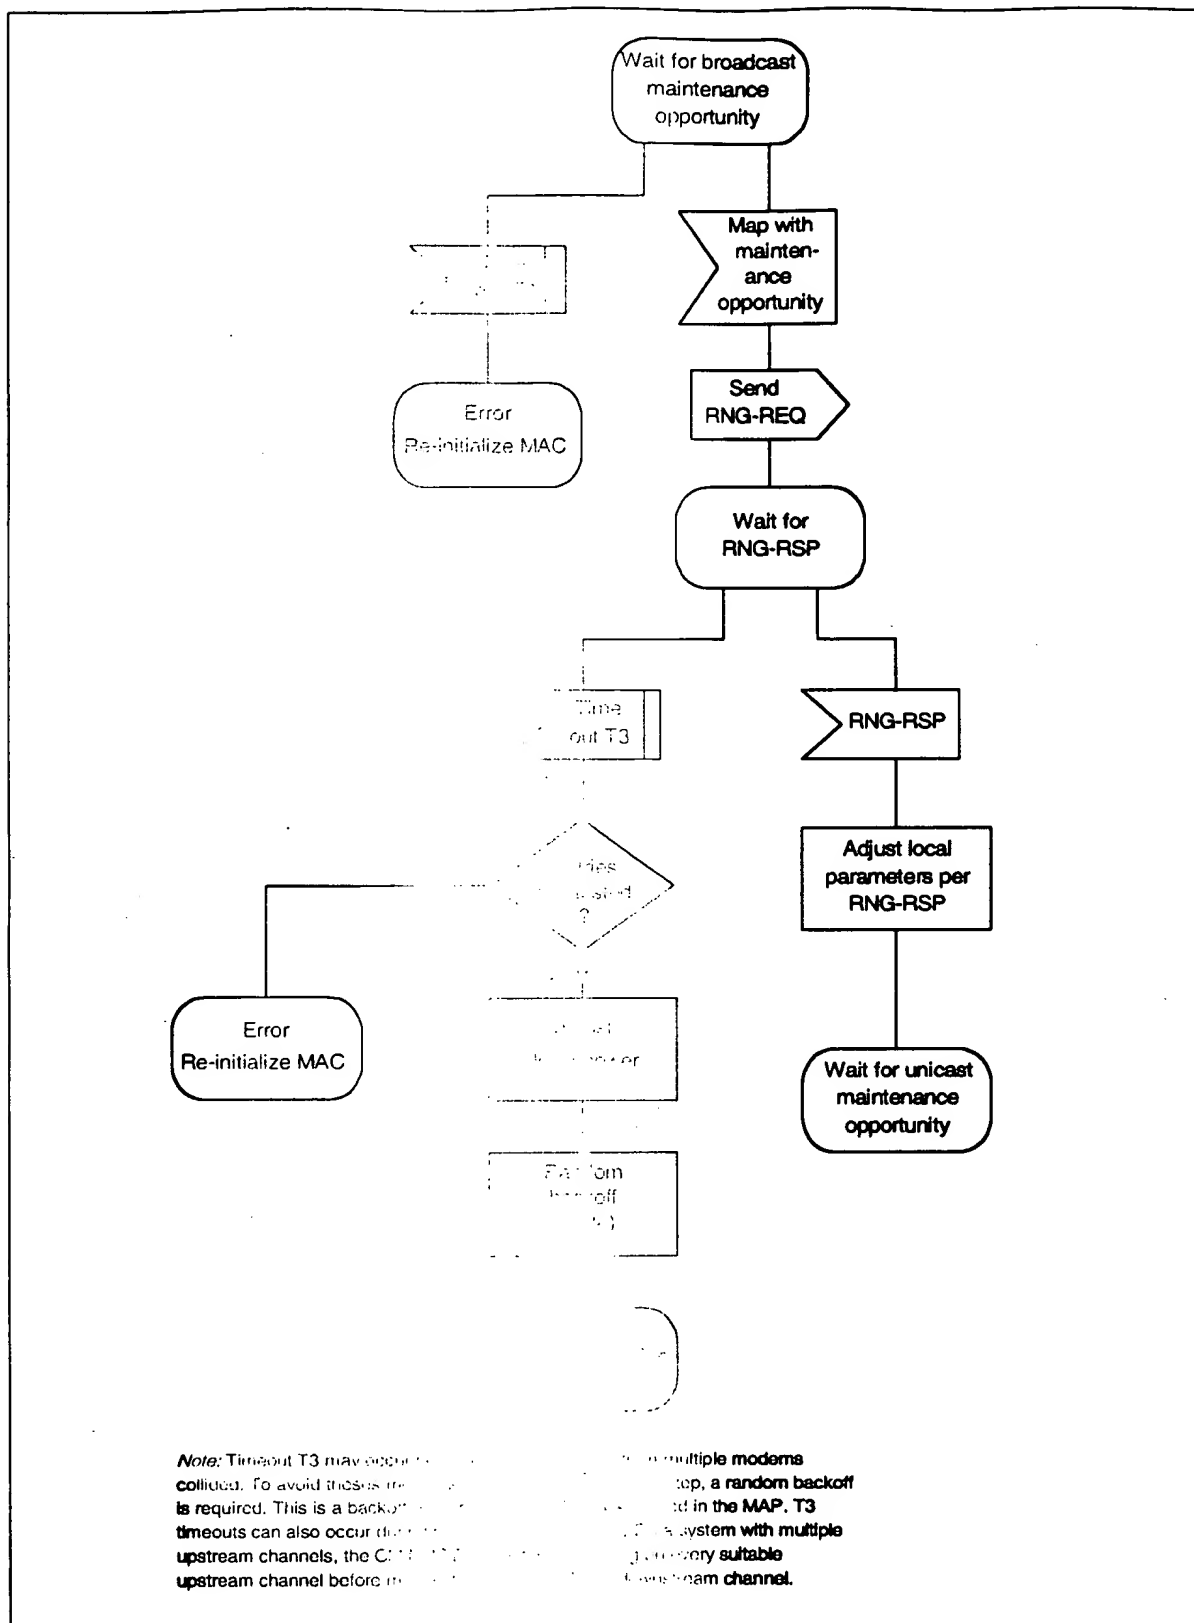
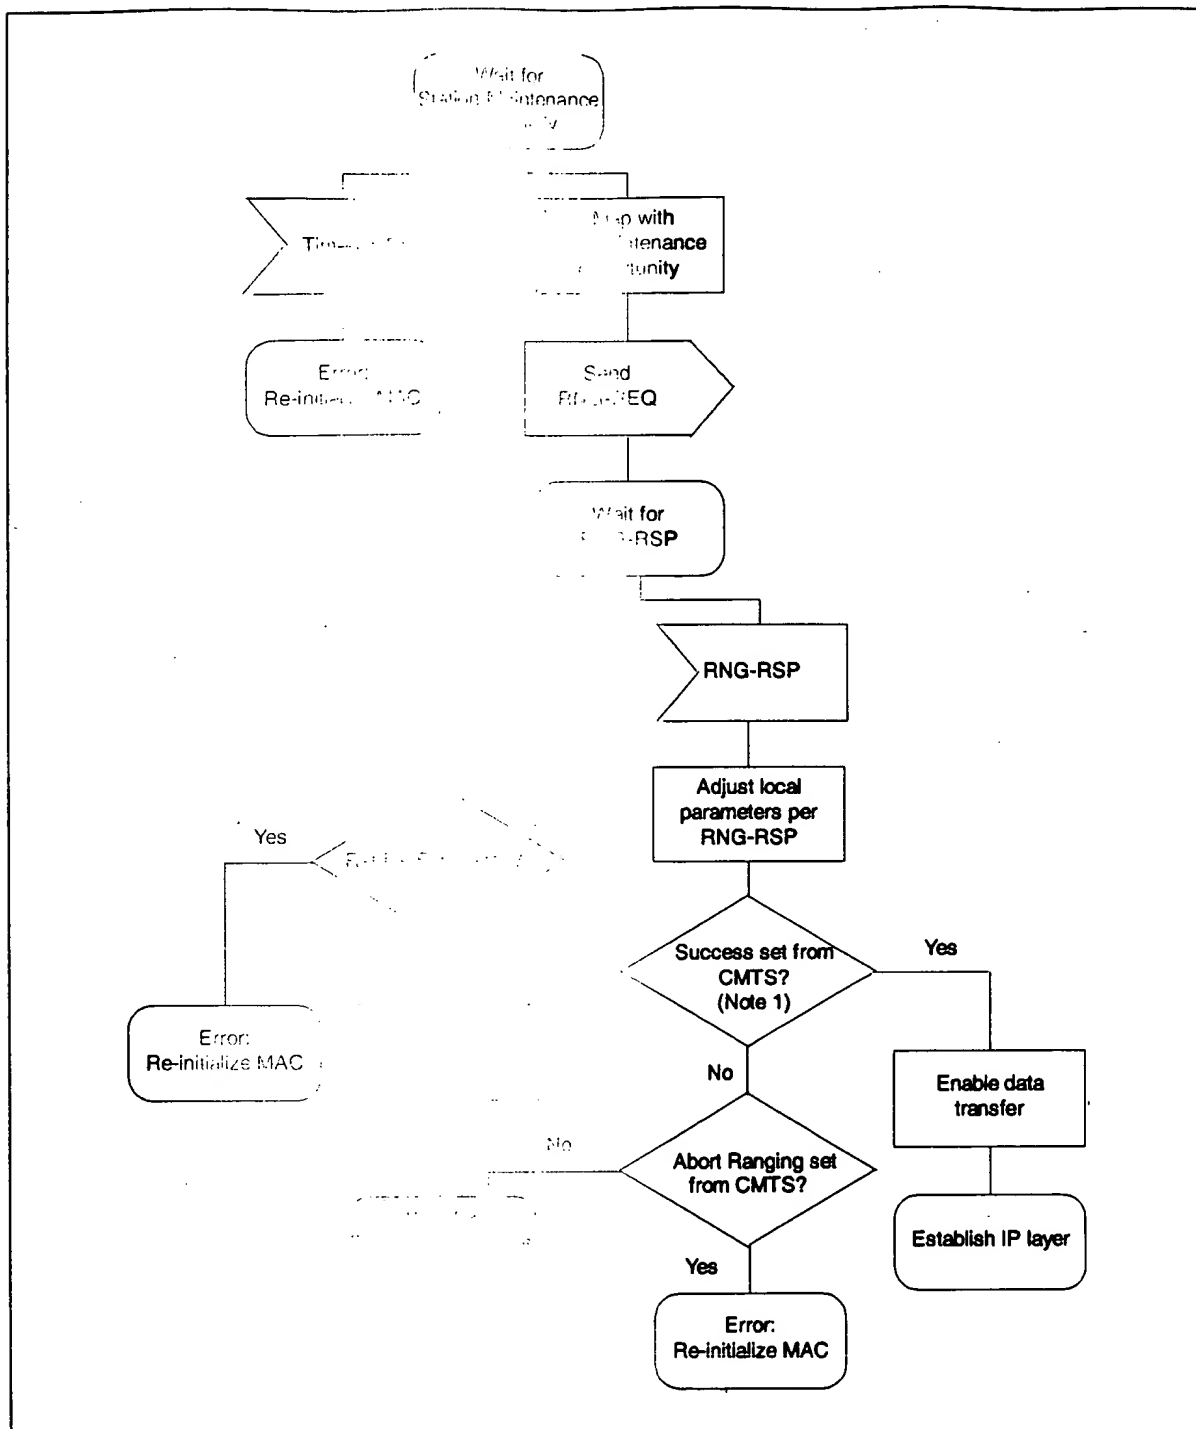


Figure 9-8: Ranging and Automatic Adjustments Procedure

Note: The CMTS **MUST** allow the CM to process the previous RNG-RSP (i.e., to modify the transmitter parameters) before sending the CMTS Ranging Response. This is defined as CM Ranging Response Time in Appendix B.

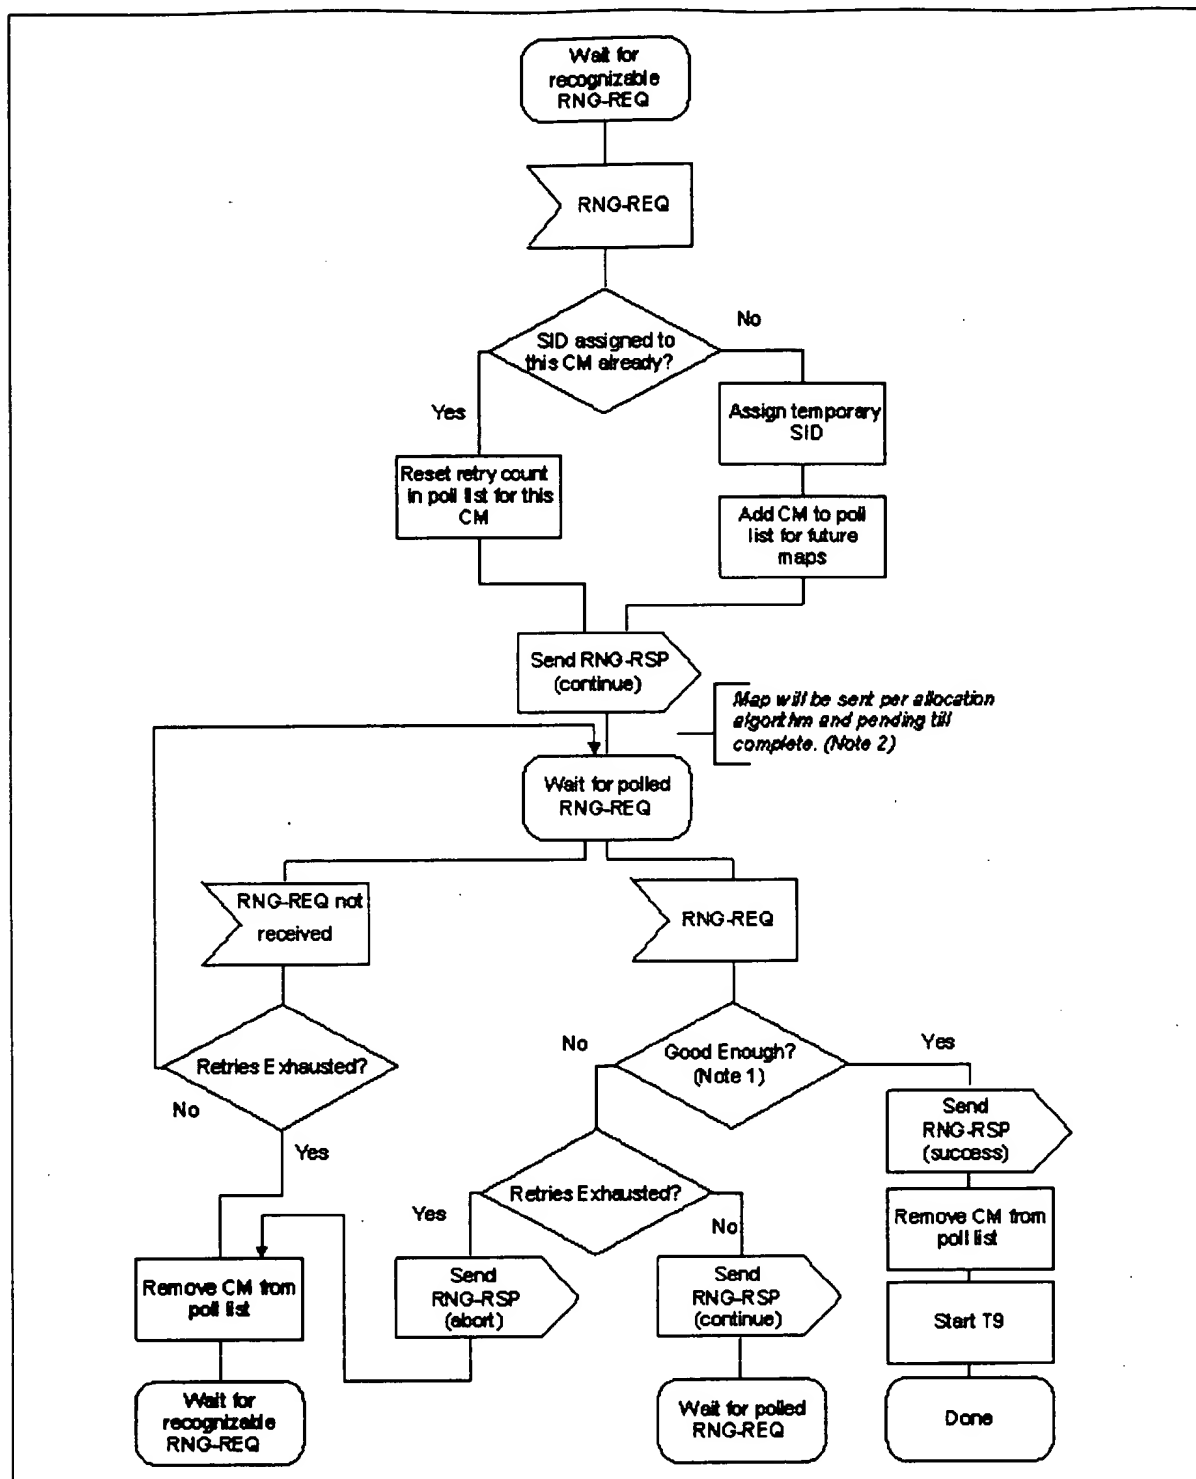


ing - CM



1. Ranging Request is received

2. Ranging - CM (continued)



1. Means ranging is within the tolerable limits of the CMTS.
2. RNG-REQ pending-till-complete was nonzero, the CMTS SHOULD hold off the station maintenance opportunity accordingly unless needed, for example, to adjust the CM's power level. If opportunities are offered prior to the pending-till-complete expiry, the "good-enough" test which follows receipt of a RNG-RSP MUST NOT judge the CM's transmit equalization until pending-till-complete expires.

Figure 9-8. Initial Ranging - CMTS (fig. edited per rfi-n-99054 06/29/99. ew)

9.2.4.1 Ranging Parameter Adjustment

Adjustment of local parameters (e.g., transmit power) in a CM as a result of the receipt (or non-receipt) of an RNG-RSP is considered to be implementation-dependent with the following restrictions (refer to Section 6.3.6):

- All parameters **MUST** be within the approved range at all times
- Power adjustment **MUST** start from the minimum value unless a valid power is available from non-volatile storage, in which case this **MUST** be used as a starting point.
- Power adjustment **MUST** be capable of being reduced or increased by the specified amount in response to RNG-RSP messages.
- If, during initialization, power is increased to the maximum value (without a response from the CMTS) it **MUST** wrap back to the minimum.
- For multi-channel support, the CM **MUST** attempt initial ranging on every suitable upstream channel before moving to the next available downstream channel.

9.2.5 Establish IP Connectivity

At this point, the CM **MUST** invoke DHCP mechanisms [RFC-2131] in order to obtain an IP address and any other parameters needed to establish IP connectivity (refer to Appendix D). The DHCP response **MUST** contain the name of a file which contains further configuration parameters. Refer to Figure 9-9.

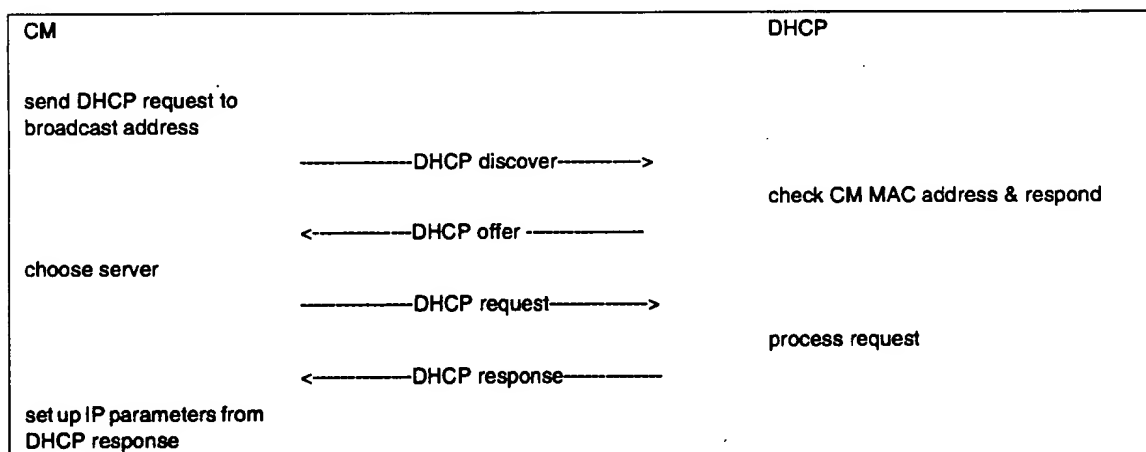


Figure 9-9. Establishing IP Connectivity

9.2.6 Establish Time of Day

The CM and CMTS need to have the current date and time. This is required for time-stamping logged events which can be retrieved by the management system. This need not be authenticated and need only be accurate to the nearest second.

The protocol by which the time of day **MUST** be retrieved is defined in [RFC-868]. Refer to Figure 9-10. The request and response **MUST** be transferred using UDP. The time retrieved from the server (UTC) **MUST** be combined with the time offset received from the DHCP response to create the current local time.

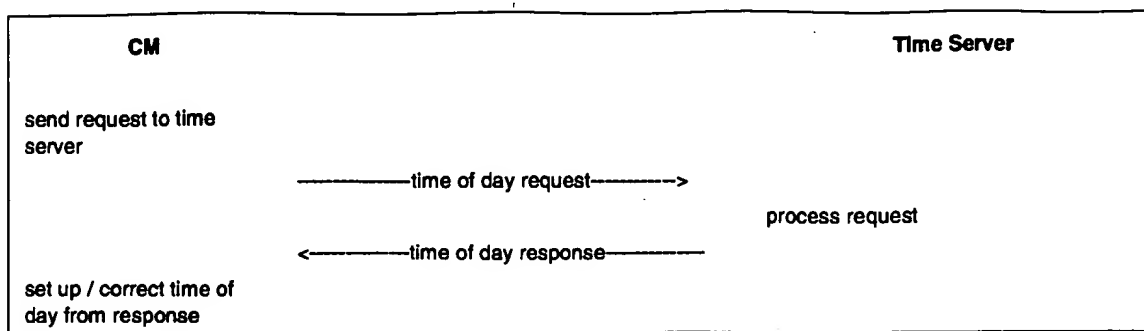


Figure 9-10. Establishing Time of Day

Successfully acquiring the Time of Day is not mandatory for a successful registration, but is necessary for on-going operation. The specific timeout for Time of Day Requests is implementation dependent. However, the CM **MUST NOT** exceed more than 3 Time of Day requests in any 5 minute period.

9.2.7 Transfer Operational Parameters

After DHCP is successful, the modem **MUST** download the parameter file using TFTP, as shown in Figure 9-11. The TFTP configuration parameter server is specified by the "siaddr" field of the DHCP response. The CM **MUST** use an adaptive timeout for TFTP based on binary exponential backoff. Refer to [RFC1123] and [RFC2349].

The parameter fields required in the DHCP response and the format and content of the configuration file **MUST** be as defined in Appendix C. Note that these fields are the minimum required for interoperability.

If a modem downloads a configuration file containing an upstream channel and/or downstream frequency different from what the modem is currently using, the modem **MUST NOT** send a Registration Request message to the CMTS. The modem **MUST** redo initial ranging using the configured upstream channel and/or downstream frequency per Section 6.3.6.3.

9.2.8 Registration

A CM **MUST** be authorized to forward traffic into the network once it is initialized and configured. The CM is authorized to forward traffic into the network via registration. To register with a CMTS, the CM **MUST** forward its configured class of service and any other operational parameters in the configuration file (refer to Section 6.3.7) to the CMTS as part of a Registration Request. Figure 9-11 shows the procedure that **MUST** be followed by the CM.

The configuration parameters downloaded to the CM **MUST** include a network access control object (see Section C.1.1.3). If this is set to "no forwarding," the CM **MUST NOT** forward data from attached CPE to the network, yet the CM **MUST** respond to network management requests. This allows the CM to be configured in a mode in which it is manageable but will not forward data.¹

1. Paragraph added per rfi-n-99052 06/29/99. ew

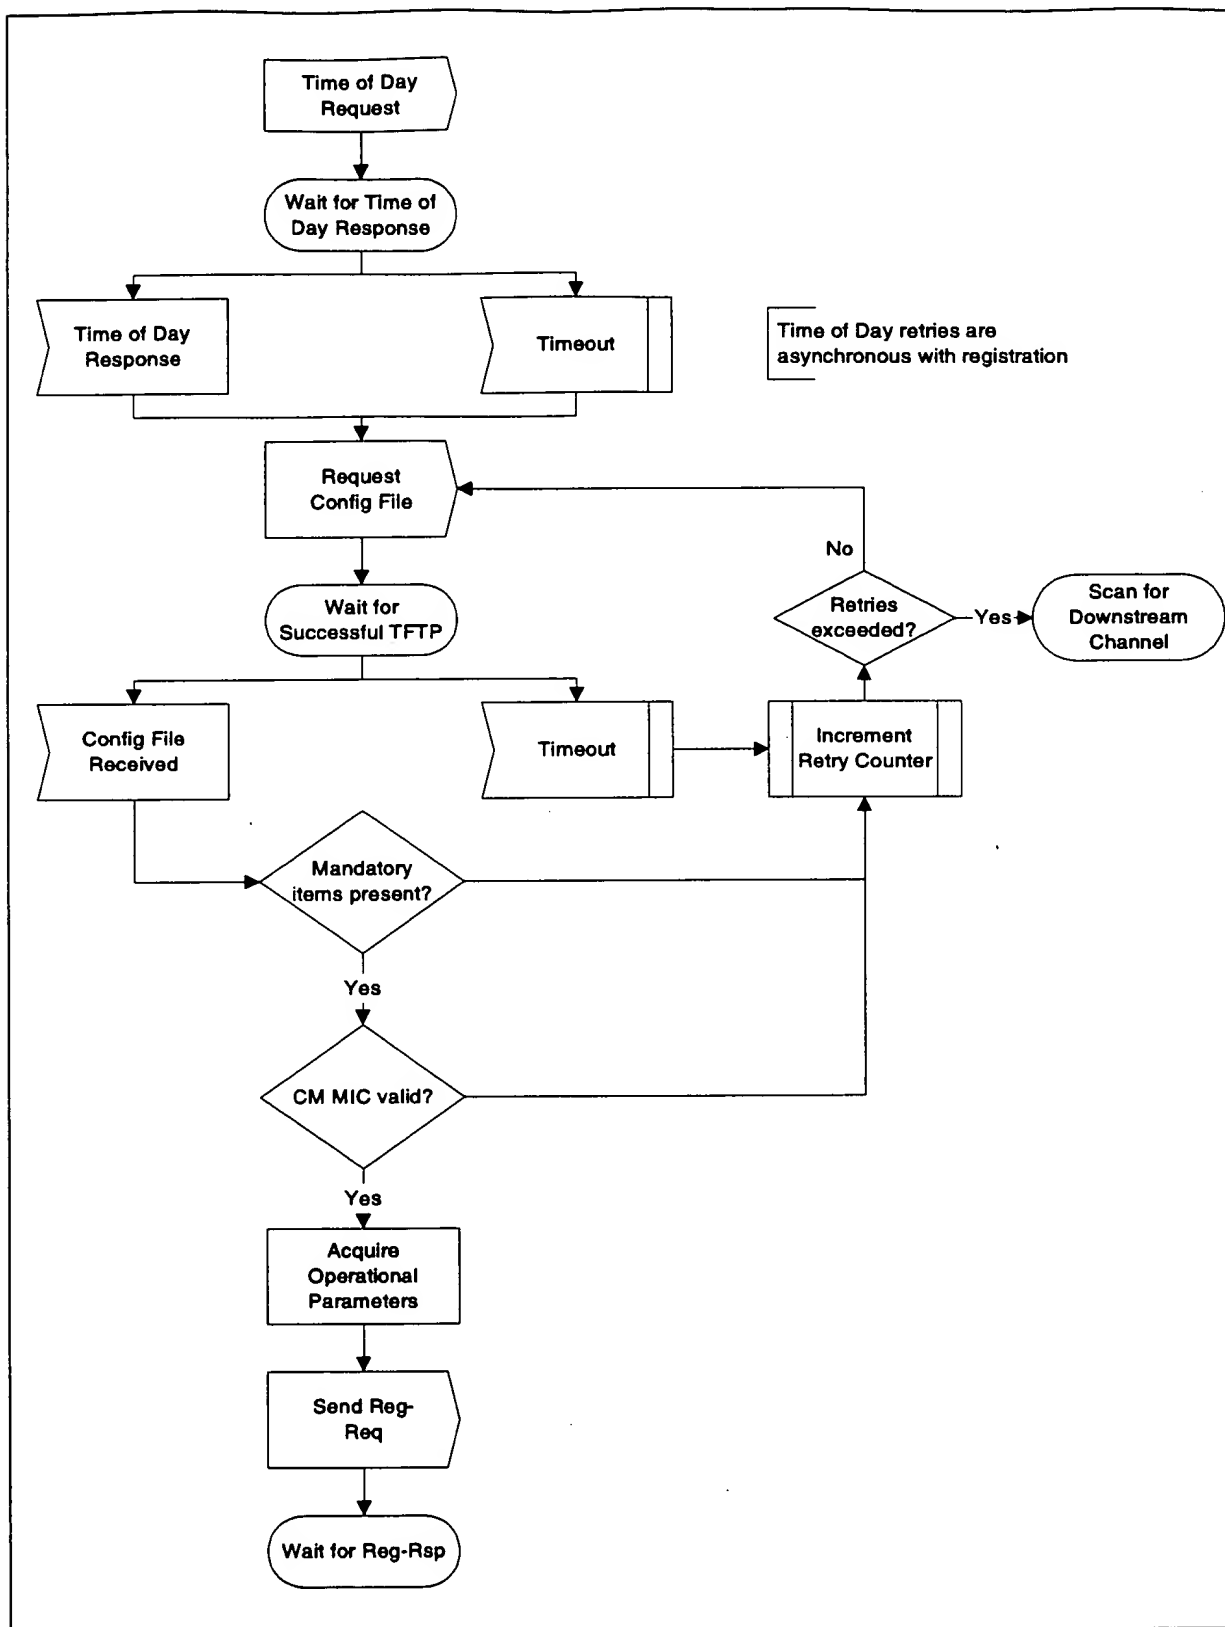


Figure 9-11. Registration — CM

Once the CM has sent a Registration Request to the CMTS it **MUST** wait for a Registration Response to authorize it to forward traffic to the network. Figure 9-12 shows the waiting procedure that **MUST** be followed by the CM.

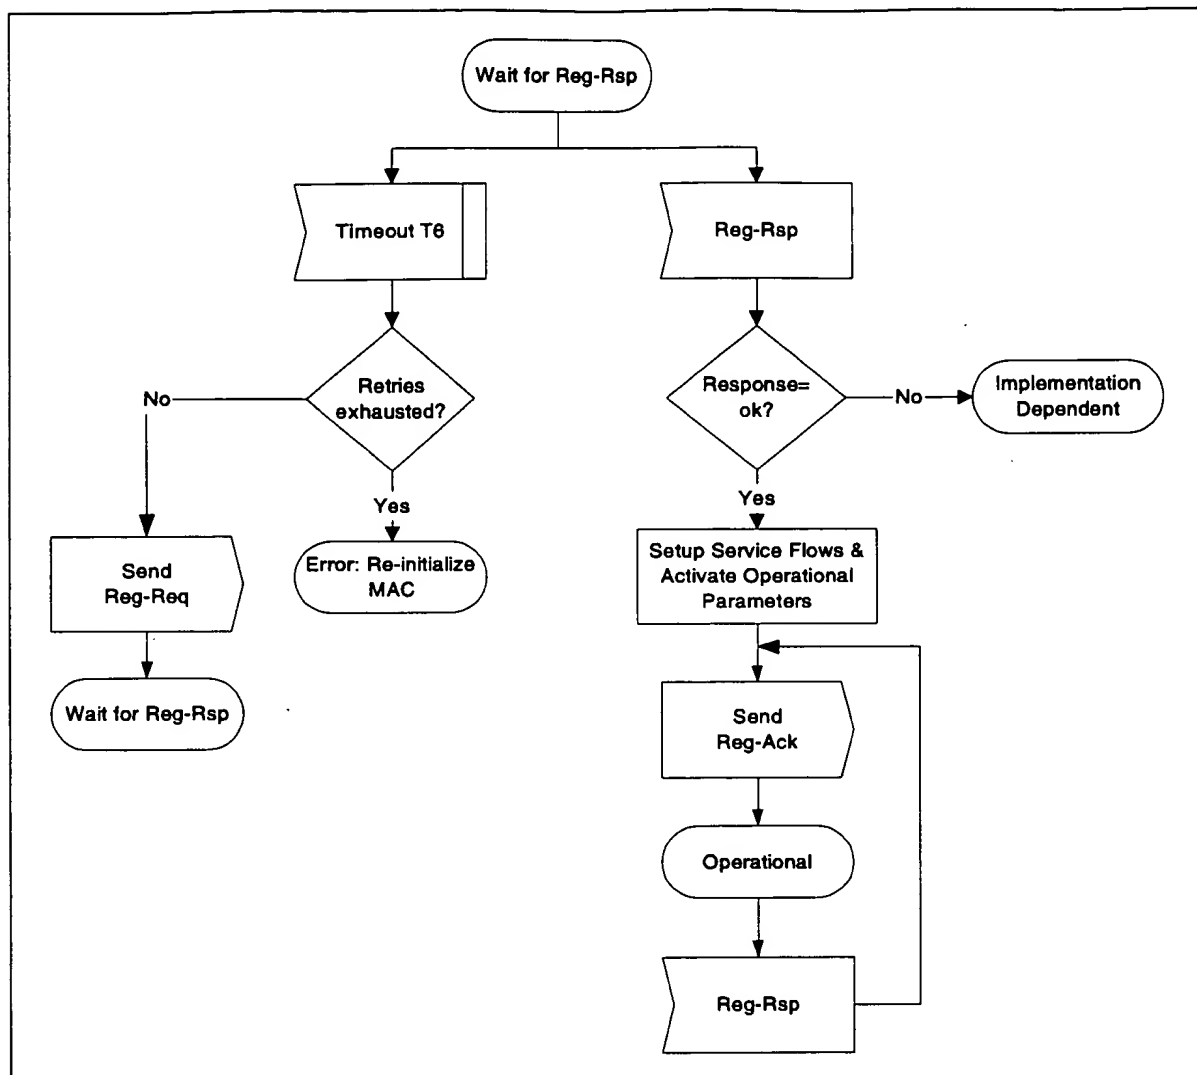


Figure 9-12. Wait for Registration Response — CM

The CMTS **MUST** perform the following operations to confirm the CM authorization (refer to Figure 9-13):

- Calculate a MIC per D.3.1 and compare it to the CMTS MIC included in the Registration Request. If the MIC is invalid, the CMTS **MUST** respond with an Authorization Failure.
- If present, check the TFTP Server Timestamp field. If the CMTS detects that the time is different from its local time by more than CM Configuration Processing Time (refer to Appendix B), the CMTS **MUST** indicate authentication failure in the REG-RSP. The CMTS **SHOULD** also make a log entry stating the CM MAC address from the message.
- If present, check the TFTP Server Provisioned Modem Address field. If the Provisioned Modem Address does not match the requesting modem's actual address, the CMTS **MUST** indicate authentication failure in the REG-RSP. The CMTS **SHOULD** also make a log entry stating the CM MAC address from the message.

- If the Registration Request contains DOCSIS 1.0 Class of Service encodings, verify the availability of the class(es) of service requested. If unable to provide the class(es) of service, the CMTS MUST respond with a Class of Service Failure and the appropriate Service Not Available response code(s). (refer to C.1.3.4)
- If the Registration Request contains Service Flow encodings, verify the availability of the Quality of Service requested in the provisioned Service Flow(s). If unable to provide the Service Flow(s), the CMTS MUST respond with a Class of Service Failure and the appropriate Service Flow Response(s).
- If the Registration Request contains DOCSIS 1.0 Class of Service encodings and Service Flow encodings, the CMTS MUST respond with a Class of Service Failure and a Service Not Available response code set to 'reject-permanent' for all DOCSIS 1.0 Classes and Service Flows requested.
- Verify the availability of any Modem Capabilities requested. If unable or unwilling to provide the Modem Capability requested, the CMTS MUST turn that Modem Capability 'off' (refer to 6.3.8.1.1).
- Assign a Service Flow ID for each class of service supported.
- Reply to the modem in a Registration Response.
- If the Registration Request contains Service Flow encodings, the CMTS MUST wait for a Registration Acknowledgment as shown in Figure 9-14.
- If timer T9 expires, the CMTS MUST both de-assign the temporary SID from that CM and make some provision for aging out that SID.¹

1. Bullet added per rfi-n-99054 06/30/99. ew

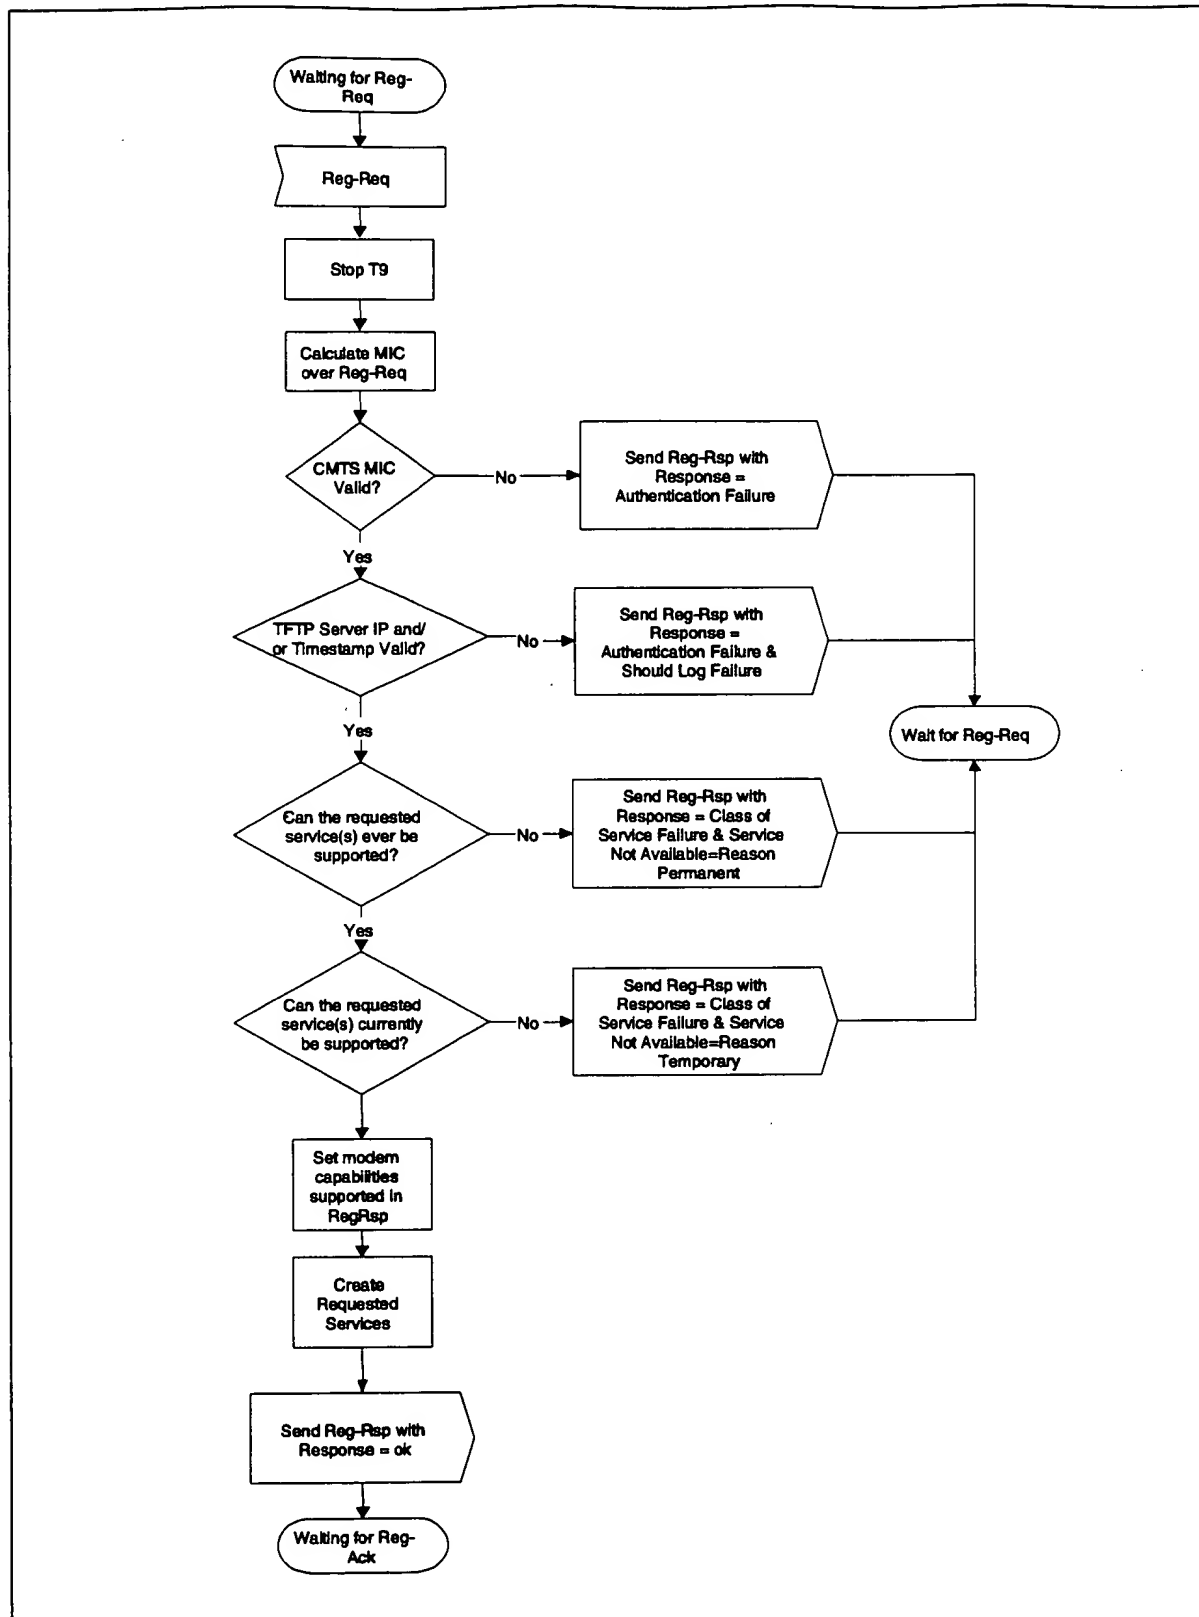


Figure 9-13. Registration — CMTS (Figure edited per rfi-n-99054 06/30/99.ew)

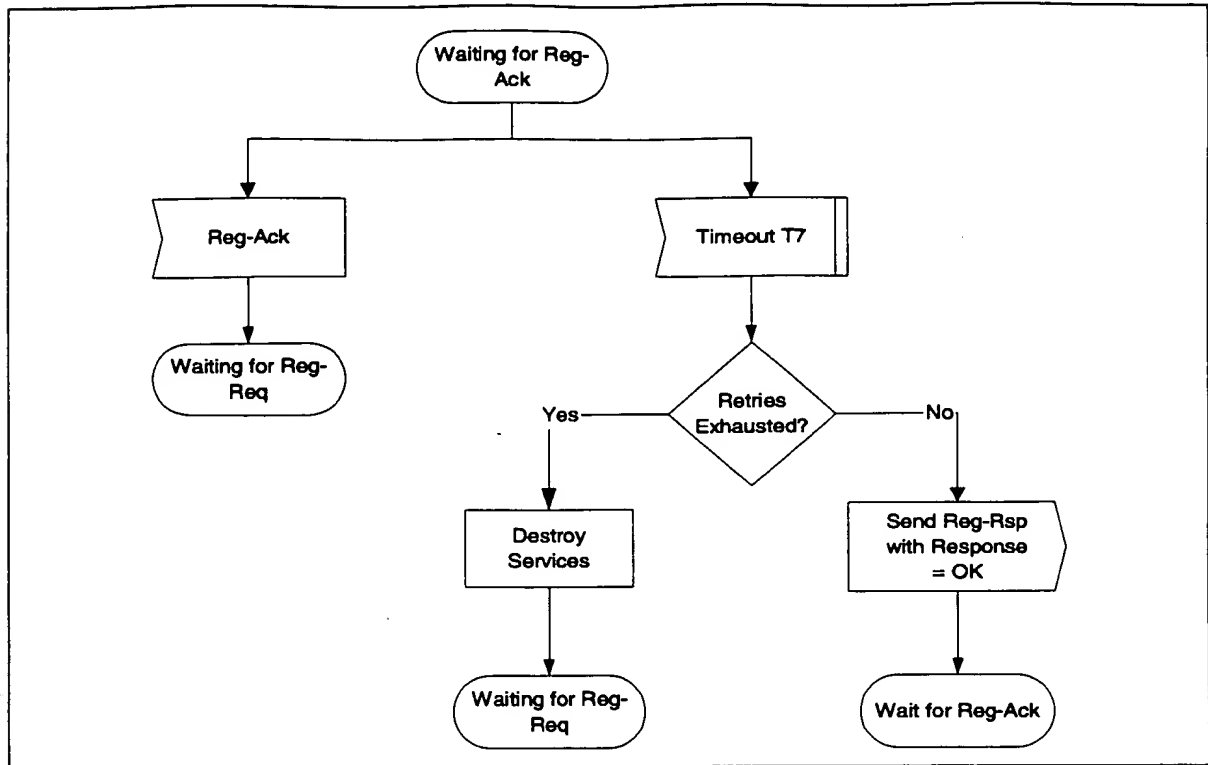


Figure 9-14. Registration Acknowledgment— CMTS

9.2.9 Baseline Privacy Initialization

Following registration, if the CM is provisioned to run Baseline Privacy, the CM **MUST** initialize Baseline Privacy operations, as described in [DOCSIS8]. A CM is provisioned to run Baseline Privacy if its configuration file includes a Baseline Privacy Configuration Setting (section C.3.2) ¹ and if the Privacy Enable parameter (section C.1.1.16) is set to enable.²

9.2.10 Service IDs During CM Initialization

After completion of the Registration process (Section 9.2.8), the CM will have been assigned Service Flow IDs (SFIDs) to match its provisioning. However, the CM must complete a number of protocol transactions prior to that time (e.g., Ranging, DHCP, etc.), and requires a temporary Service ID in order to complete those steps.

On reception of an Initial Ranging Request, the CMTS **MUST** allocate a temporary SID and assign it to the CM for initialization use. The CMTS **MAY** monitor use of this SID and restrict traffic to that needed for initialization. It **MUST** inform the CM of this assignment in the Ranging Response.

On receiving a Ranging Response addressed to it, the CM **MUST** use the assigned temporary SID for further initialization transmission requests until the Registration Response is received.

1. clarification added 06/22/99 per rfi-n-99043

2. clarification added 06/08/99 per rfi-n-99040

On receiving a Ranging Response instruction to move to a new downstream frequency and/or upstream channel ID, the CM MUST consider any previously assigned temporary SID to be deassigned, and must obtain a new temporary SID via initial ranging.

It is possible that the Ranging Response may be lost after transmission by the CMTS. The CM MUST recover by timing out and re-issuing its Initial Ranging Request. Since the CM is uniquely identified by the source MAC address in the Ranging Request, the CMTS MAY immediately re-use the temporary SID previously assigned. If the CMTS assigns a new temporary SID, it MUST make some provision for aging out the old SID that went unused (see Section 6.3.8).

When assigning provisioned SFIDs on receiving a Registration Request, the CMTS may re-use the temporary SID, assigning it to one of the Service Flows requested. If so, it MUST continue to allow initialization messages on that SID, since the Registration Response could be lost in transit. If the CMTS assigns all-new SIDs for class-of-service provisioning, it MUST age out the temporary SID. The aging-out MUST allow sufficient time to complete the registration process in case the Registration Response is lost in transit.

9.2.11 Multiple-Channel Support

In the event that more than one downstream signal is present in the system, the CM MUST operate using the first valid downstream signal that it encounters when scanning. It will be instructed via the parameters in the configuration file (see Appendix C) to shift operation to different downstream and/or upstream frequencies if necessary.

Both upstream and downstream channels MUST be identified where required in MAC management messages using channel identifiers.

9.3 Standard Operation

9.3.1 Periodic Signal Level Adjustment

The CMTS MUST provide each CM a Periodic Ranging opportunity at least once every T4 seconds. The CMTS MUST send out Periodic Ranging opportunities at an interval sufficiently shorter than T4 that a MAP could be missed without the CM timing out. The size of this "subinterval" is CMTS dependent.

The CM MUST reinitialize its MAC layer after T4 seconds have elapsed without receiving a Periodic Ranging opportunity.

Remote RF signal level adjustment at the CM is performed through a periodic maintenance function using the RNG-REQ and RNG-RSP MAC messages. This is similar to initial ranging and is shown in Figure 9-15 and Figure 9-16. On receiving a RNG-RSP, the CM MUST NOT transmit until the RF signal has been adjusted in accordance with the RNG-RSP and has stabilized (refer to Section 4).

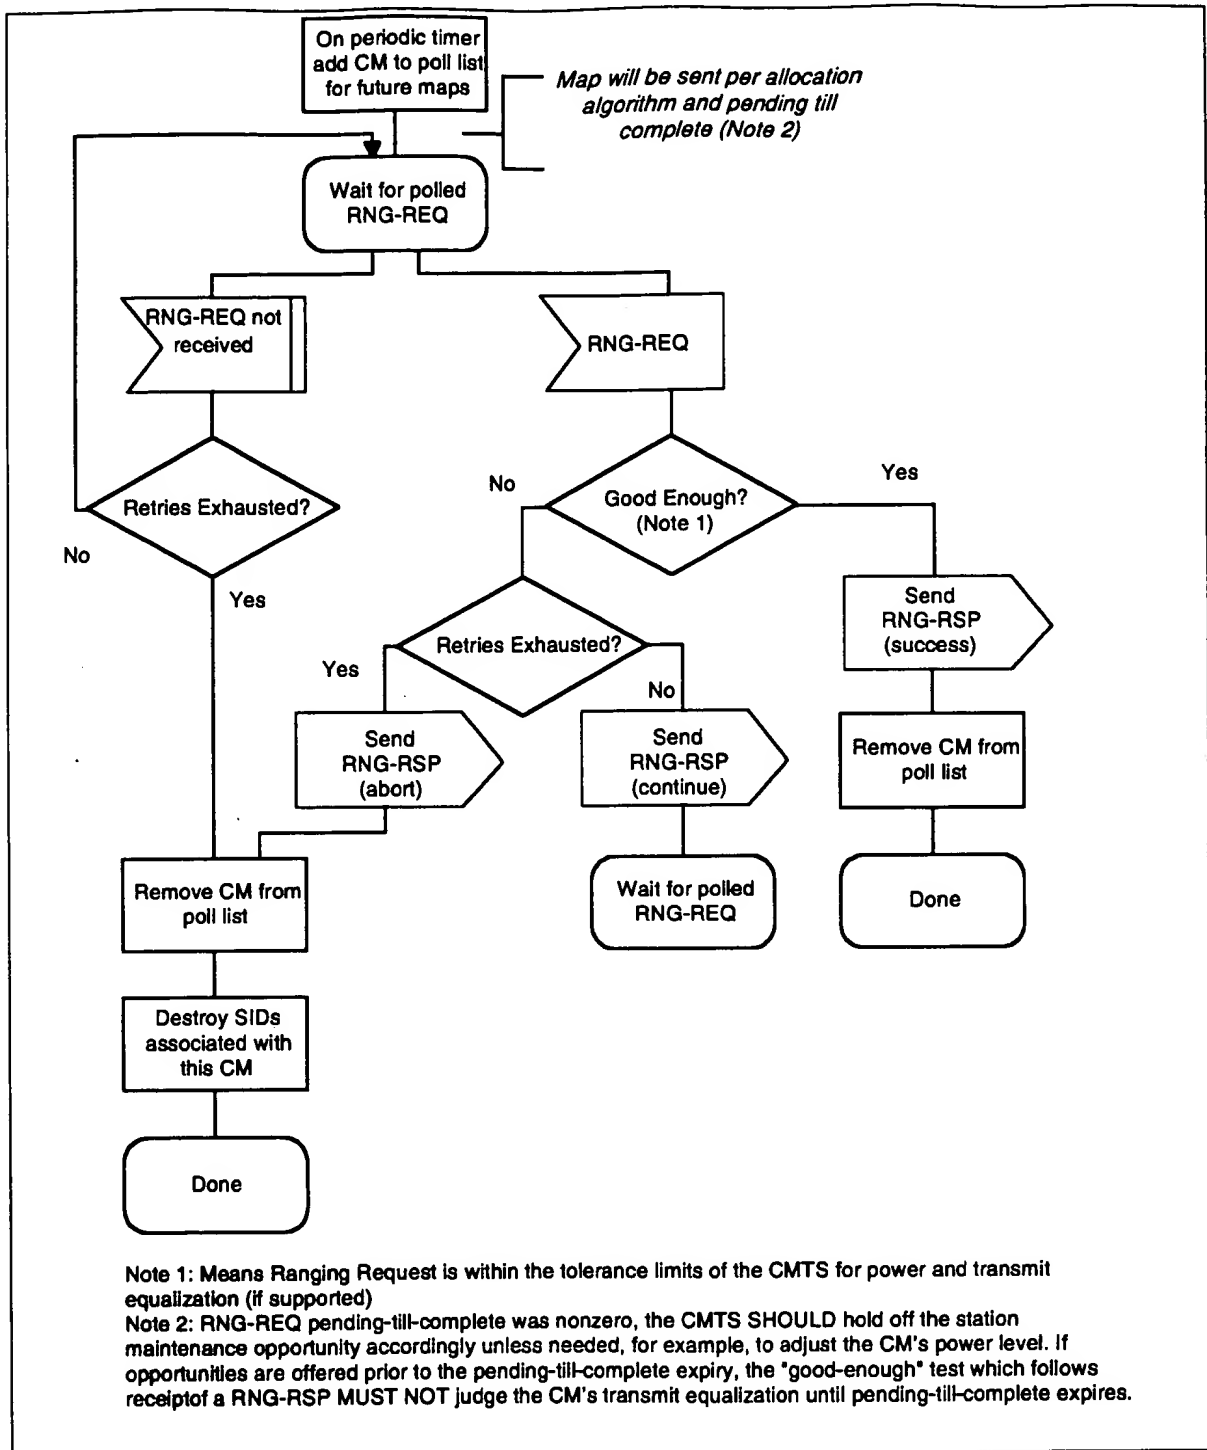


Figure 9-15. Periodic Ranging - CMTS

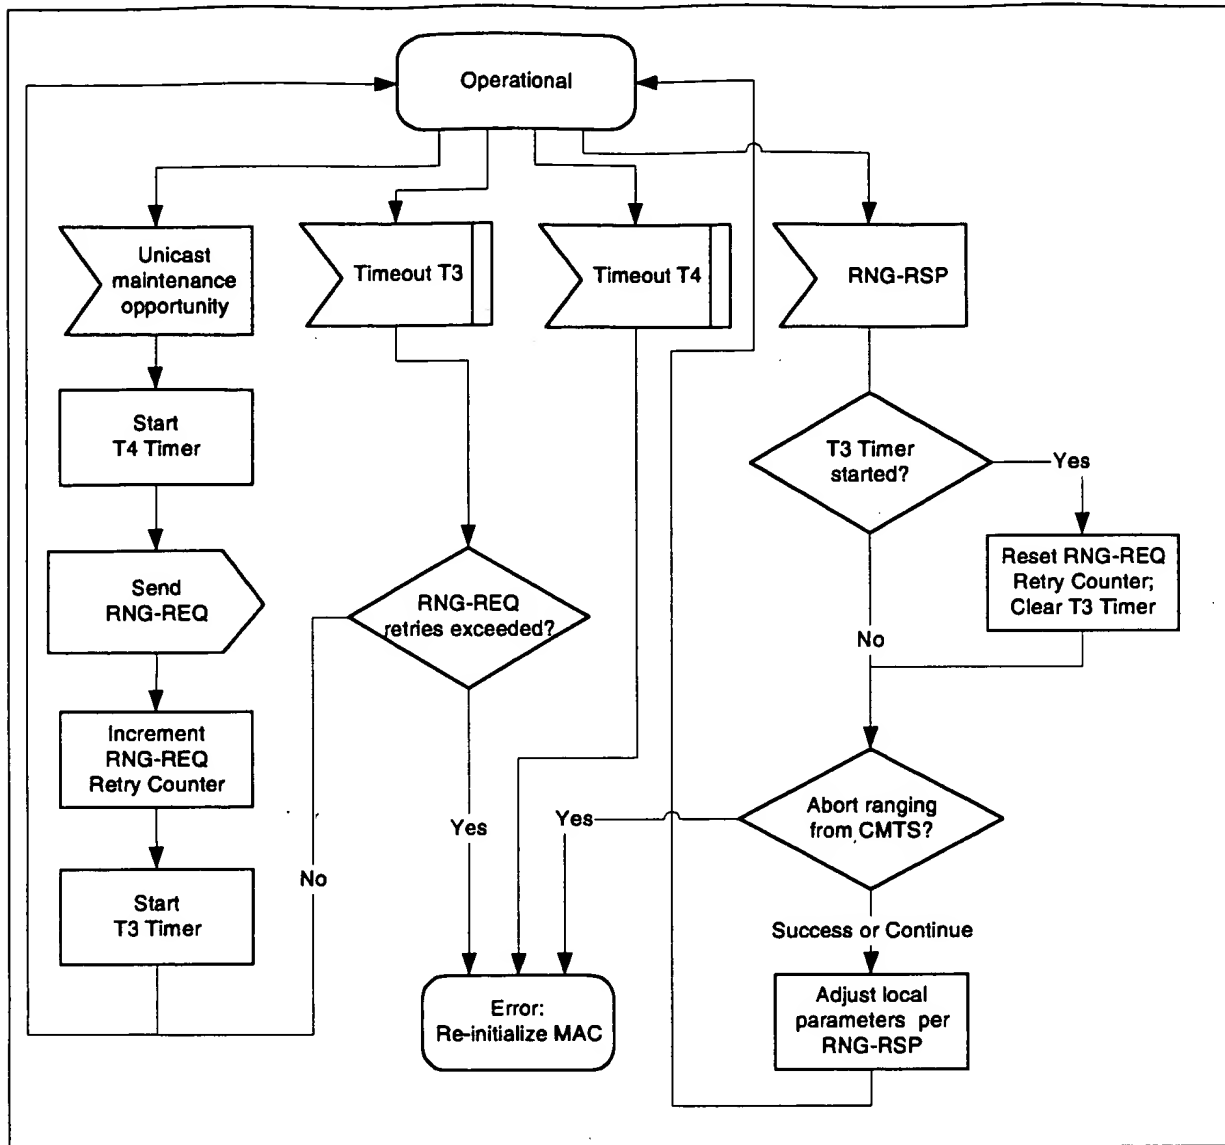


Figure 9-16. Periodic Ranging - CM View

9.3.2 Changing Upstream Burst Parameters

Whenever the CMTS is to change any of the upstream burst characteristics, it must provide for an orderly transition from the old values to the new values by all CMs. Whenever the CMTS is to change any of the upstream burst values, it **MUST**:

- Announce the new values in an Upstream Channel Descriptor message. The Configuration Change Count field must be incremented to indicate that a value has changed.

After transmitting one or more UCD messages with the new value, the CMTS transmits a MAP message with a UCD Count matching the new Configuration Change Count. The first interval in the MAP **MUST** be a data grant of at least 1 millisecond to the null Service ID (zero). That is, the CMTS **MUST** allow one millisecond for cable modems to change their PMD sublayer parameters to match the new set. This millisecond is in addition to other MAP timing constraints (see Section 7.1.5).

- The CMTS MUST NOT transmit MAPs with the old UCD Count after transmitting the new UCD.

The CM MUST use the parameters from the UCD corresponding to the MAP's "UCD Count" for any transmissions it makes in response to that MAP. If the CM has, for any reason, not received the corresponding UCD, it cannot transmit during the interval described by that MAP.

9.3.3 Changing Upstream Channels

At any time after registration, the CMTS MAY direct the CM to change its upstream channel. This may be done for traffic balancing, noise avoidance, or any of a number of other reasons which are beyond the scope of this specification. Figure 9-17 shows the procedure that MUST be followed by the CMTS. Figure 9-18 shows the corresponding procedure at the CM.

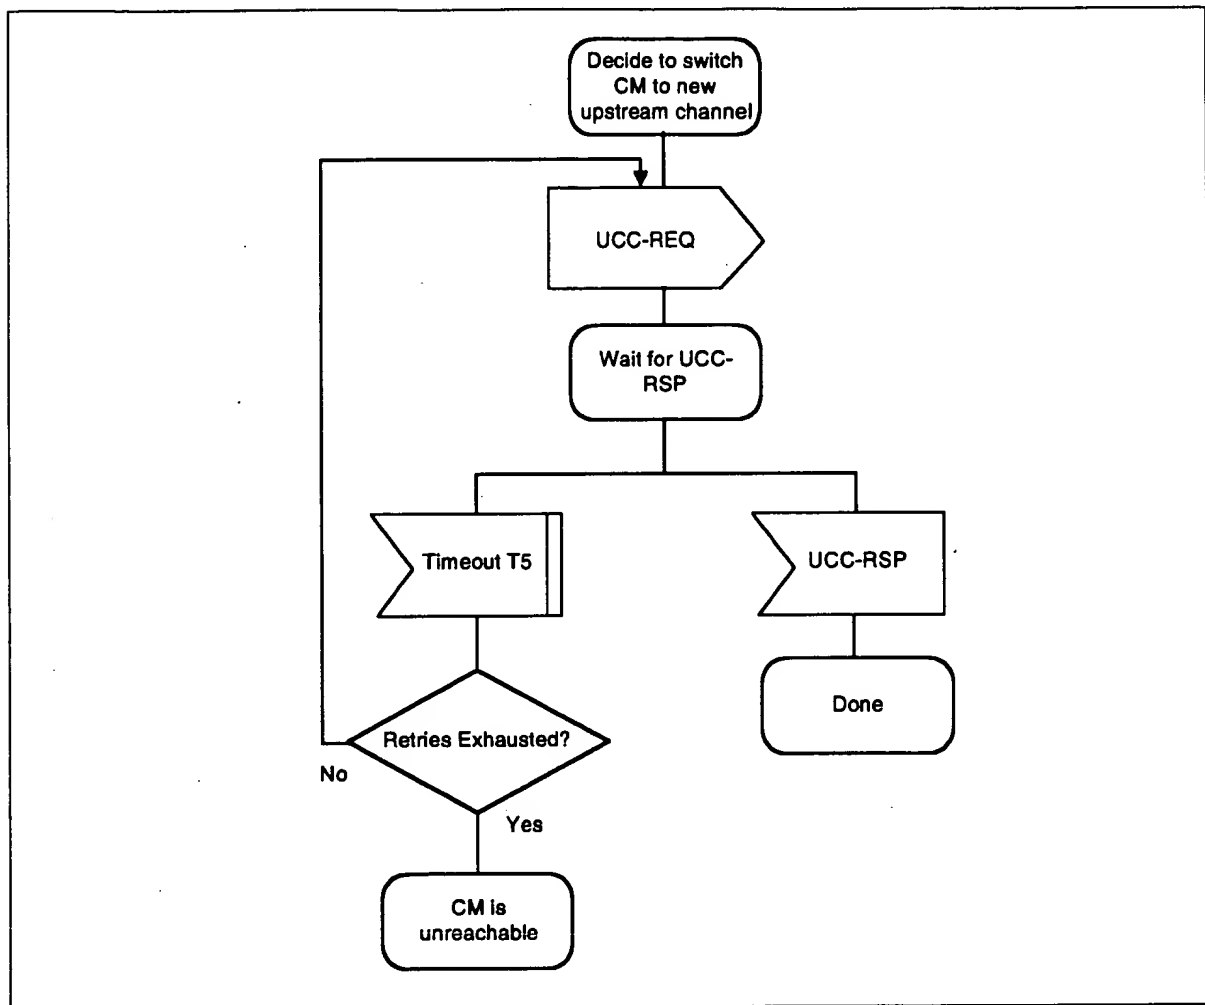


Figure 9-17. Changing Upstream Channels: CMTS View

Note that if the CMTS retries the UCC-REQ, the CM may have already changed channels (if the UCC-RSP was lost in transit). Consequently, the CMTS MUST listen for the UCC-RSP on both the old and the new channels.

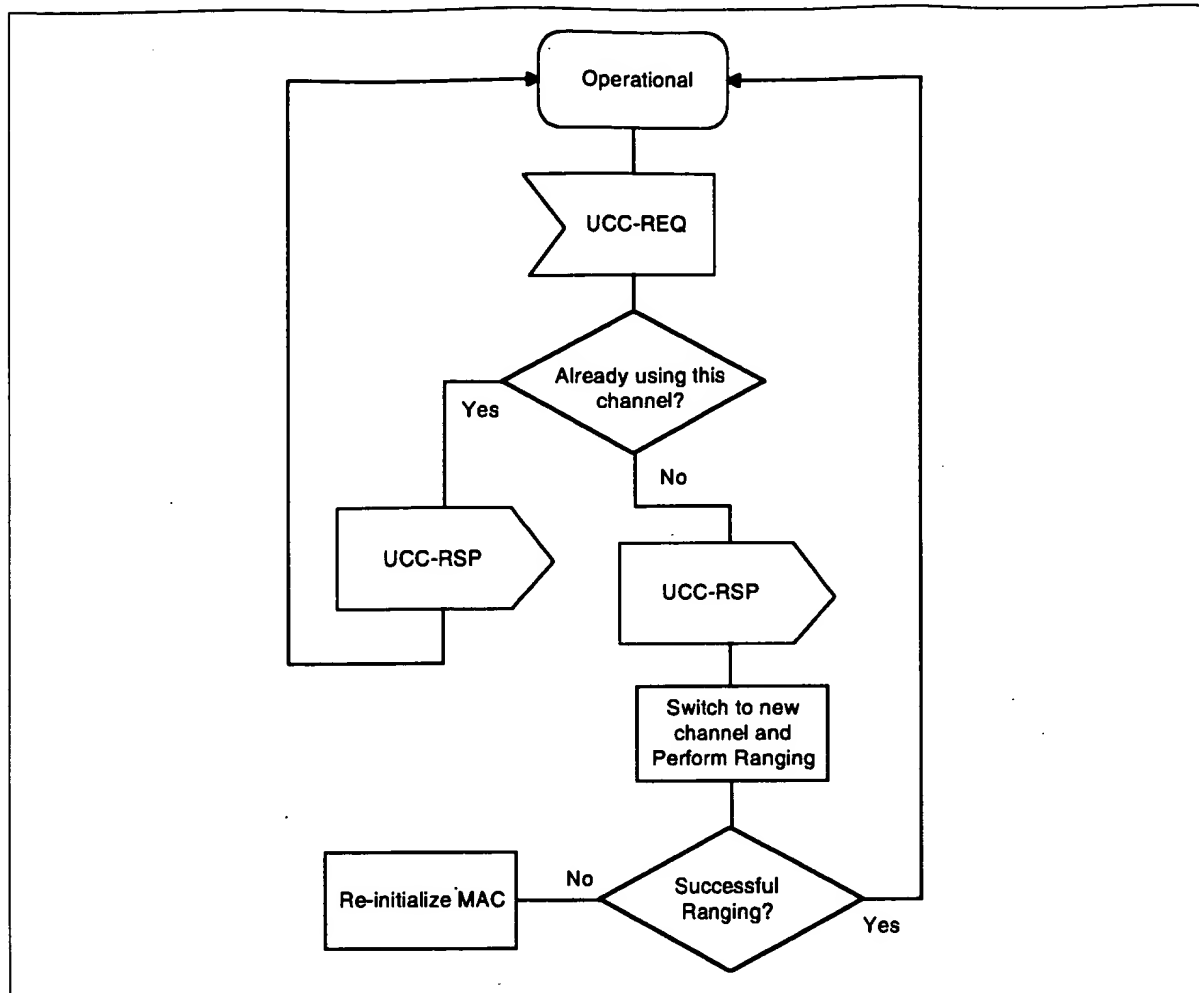


Figure 9-18. Changing Upstream Channels: CM View

Upon synchronizing with the new upstream channel, the CM **MUST** re-range using the technique specified in the UCC-REQ Ranging Technique TLV, if present. If this TLV is not present in the UCC-REQ, the CM **MUST** perform initial maintenance on the new upstream channel. (Refer to 6.3.10.1.1)

If the CM has previously established ranging on the new channel, and if that ranging on that channel is still current (T4 has not elapsed since the last successful ranging), then the CM **MAY** use cached ranging information and omit ranging.

The CM **SHOULD** cache UCD information from multiple upstream channels to eliminate waiting for a UCD corresponding to the new upstream channel.

The CM **MUST NOT** perform re-registration, since its provisioning and MAC domain remain valid on the new channel.

9.4 Dynamic Service

Service Flows MAY be created, changed or deleted. This is accomplished through a series of MAC management messages referred to as Dynamic Service Addition (DSA), Dynamic Service Change (DSC) and Dynamic Service Deletion (DSD). The DSA messages create a new Service Flow. The DSC messages change an existing Service Flow. The DSD messages delete an existing Service Flow. This is illustrated in Figure 9-19.¹




Figure 9-19. Dynamic Service Flow Overview

The START state implies that no Service Flow exists that matches the SFID and/or Transaction ID in a message. Once the Service Flow exists, it is in the SF_operational state and has an assigned SFID. Since multiple Service Flows MAY exist, there may be multiple state machines active, one for every Service Flow. Dynamic Service messages only affect those state machines that match the SFID and/or Transaction ID. If privacy is enabled, both the CM and CMTS MUST verify the HMAC digest on all dynamic service messages before processing them, and discard any messages that fail.²

Service Flows created at registration time effectively enter the SF_operational state without a DSA transaction.

TransactionIDs are unique per sender. That is, a CM originated transaction is completely unique from any CMTS originated transaction, even if the TransactionIDs are equal.

Each dynamic service message sequence is a unique transaction with an associated unique transaction identifier. The DSA/DSC transactions consist of a request/response/acknowledge sequence. The DSD transactions consist of a request/response sequence. The response messages will return a confirmation code of okay unless some exception condition was detected. The acknowledge messages will return the confirmation code in the response unless a new exception condition arises. A more detailed state diagram, including transition states, is shown below. Separate diagrams are provided for the CM and CMTS as there are subtle differences between them. The detailed actions for each transaction will be given in the following sections.

1. text added 06/22/99 per rfi-n-99043 ew

2. text added 06/22/99 per rfi-n-99043 ew

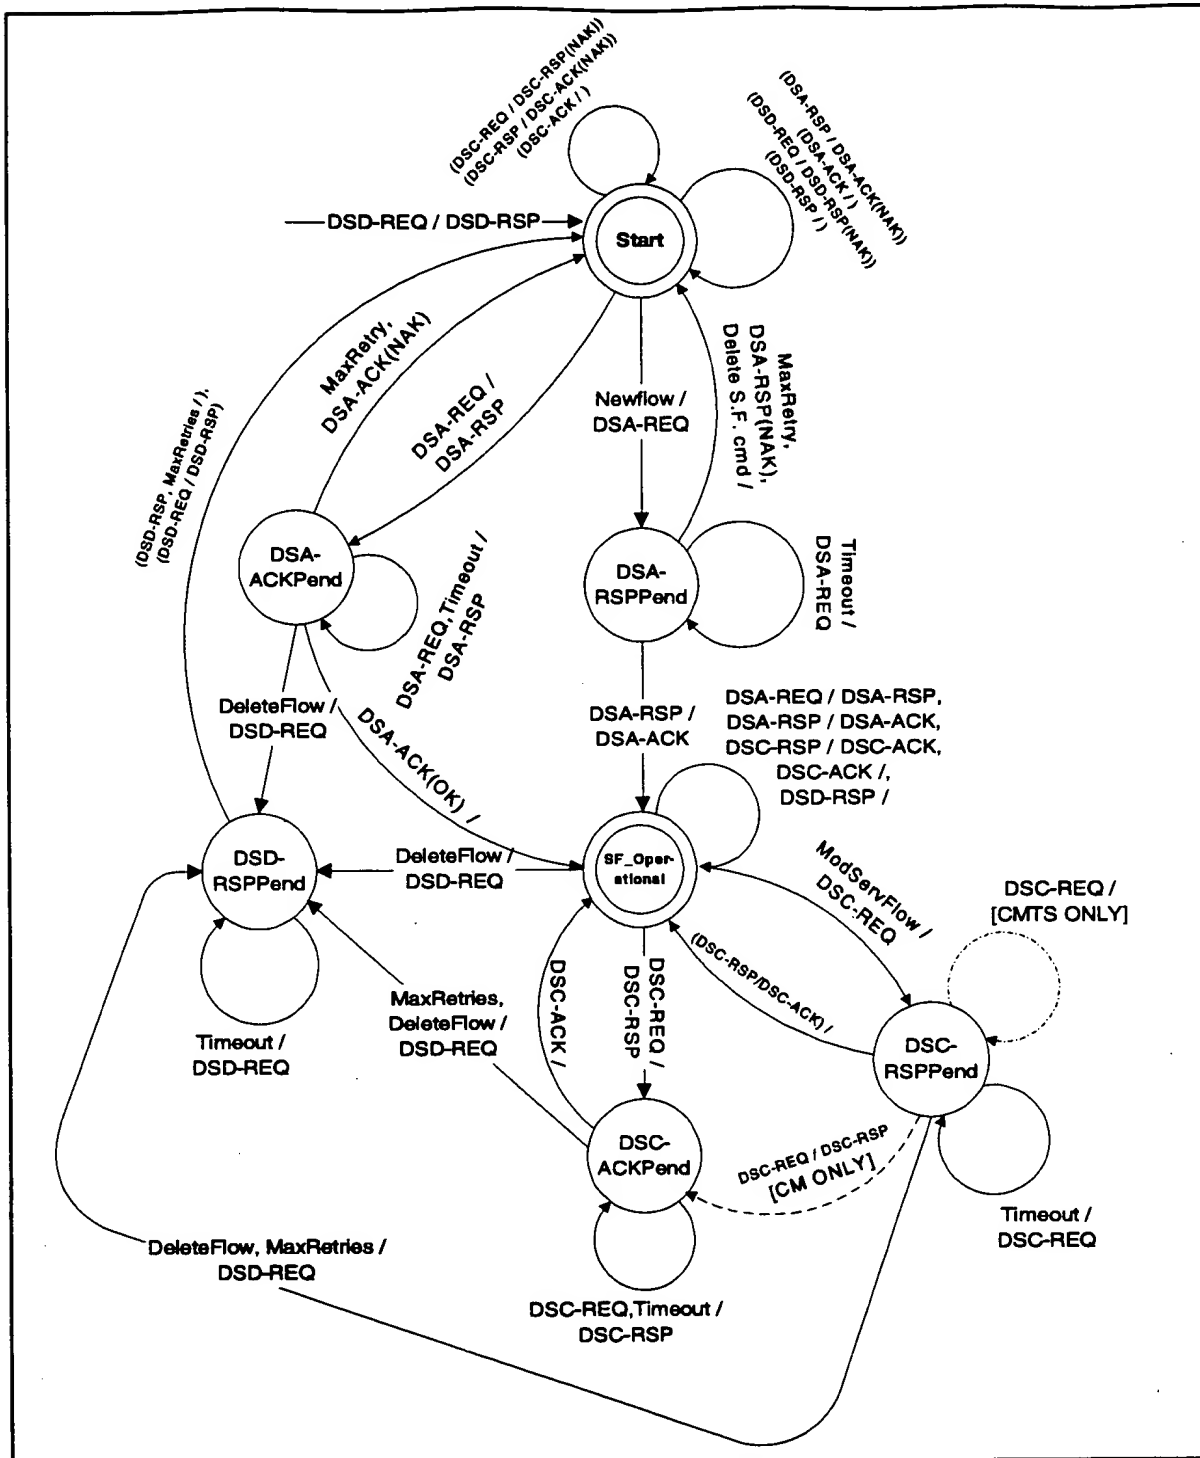


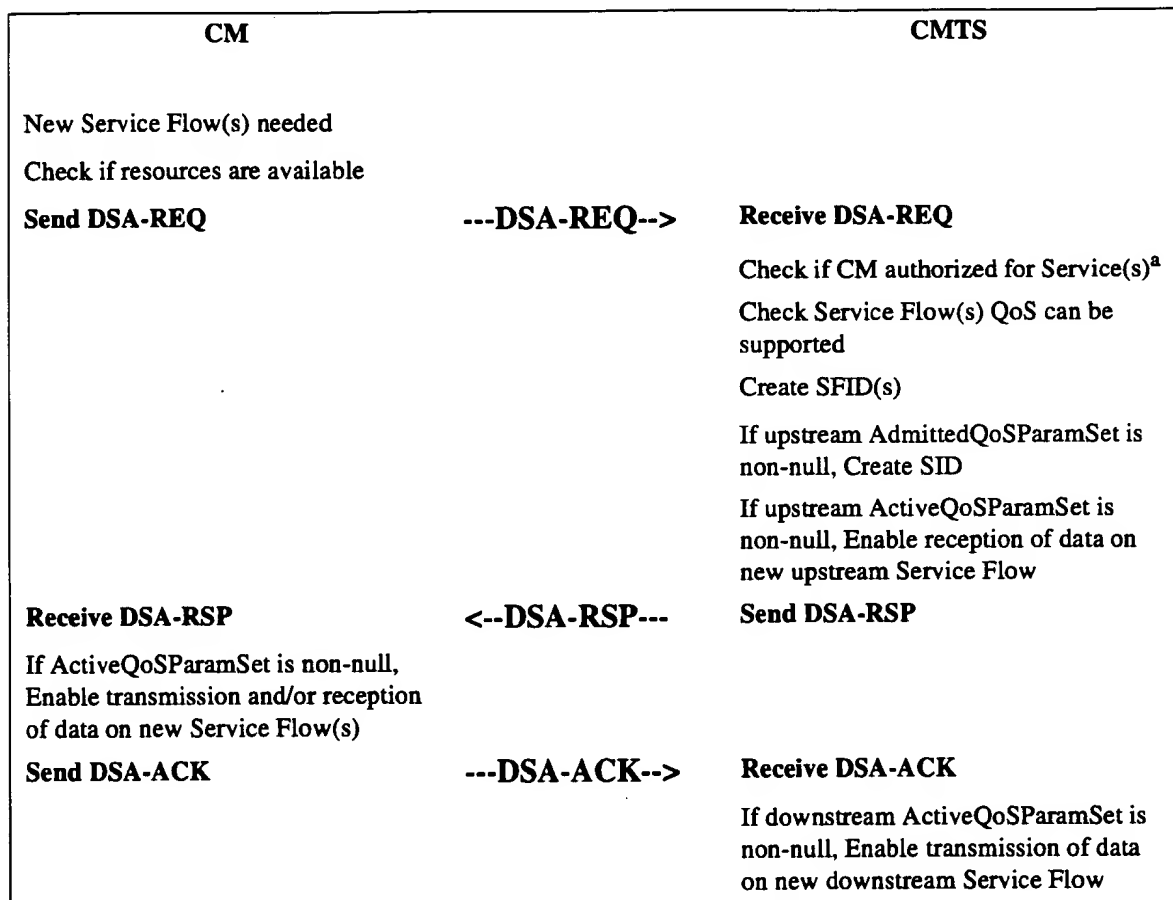
Figure 9-20. Dynamic Service Flow State Transition Diagram

9.4.1 Dynamic Service Addition

9.4.1.1 CM Initiated Dynamic Service Addition

A CM wishing to create an upstream and/or a downstream Service Flow sends a request to the CMTS using a dynamic service addition request message (DSA-REQ)¹. The CMTS checks the CM's authorization for the requested service(s) and whether the QoS requirements can be supported and generates an appropriate response using a dynamic service addition response message (DSA-RSP). The CM concludes the transaction with an acknowledgement message (DSA-ACK).

In order to facilitate a common admission response, an upstream and a downstream Service Flow can be included in a single DSA-REQ. Both Service Flows are either accepted or rejected together.²



a.Note: authorization can happen prior to the DSA-REQ being received by the CMTS. The details of CMTS signalling to anticipate a DSA-REQ are beyond the scope of this specification.

Figure 9-21. Dynamic Service Addition Initiated from CM (Figure edited per rfi-n-99048 06/30/99. ew)

1. Sentence edited per rfi-n-99048 06/30/99. ew

2. Paragraph edited per rfi-n-99048 06/30/99. ew

9.4.1.2 CMTS Initiated Dynamic Service Addition

A CMTS wishing to establish an upstream and/or a downstream dynamic Service Flow(s) with a CM performs the following operations.¹ The CMTS checks the authorization of the destination CM for the requested class of service and whether the QoS requirements can be supported. If the service can be supported the CMTS generates new SFID(s) with the required class of service and informs the CM using a dynamic service addition request message (DSA-REQ). If the CM checks that it can support the service and responds using a dynamic service addition response message (DSA-RSP). The transaction completes with the CMTS sending the acknowledge message (DSA-ACK).

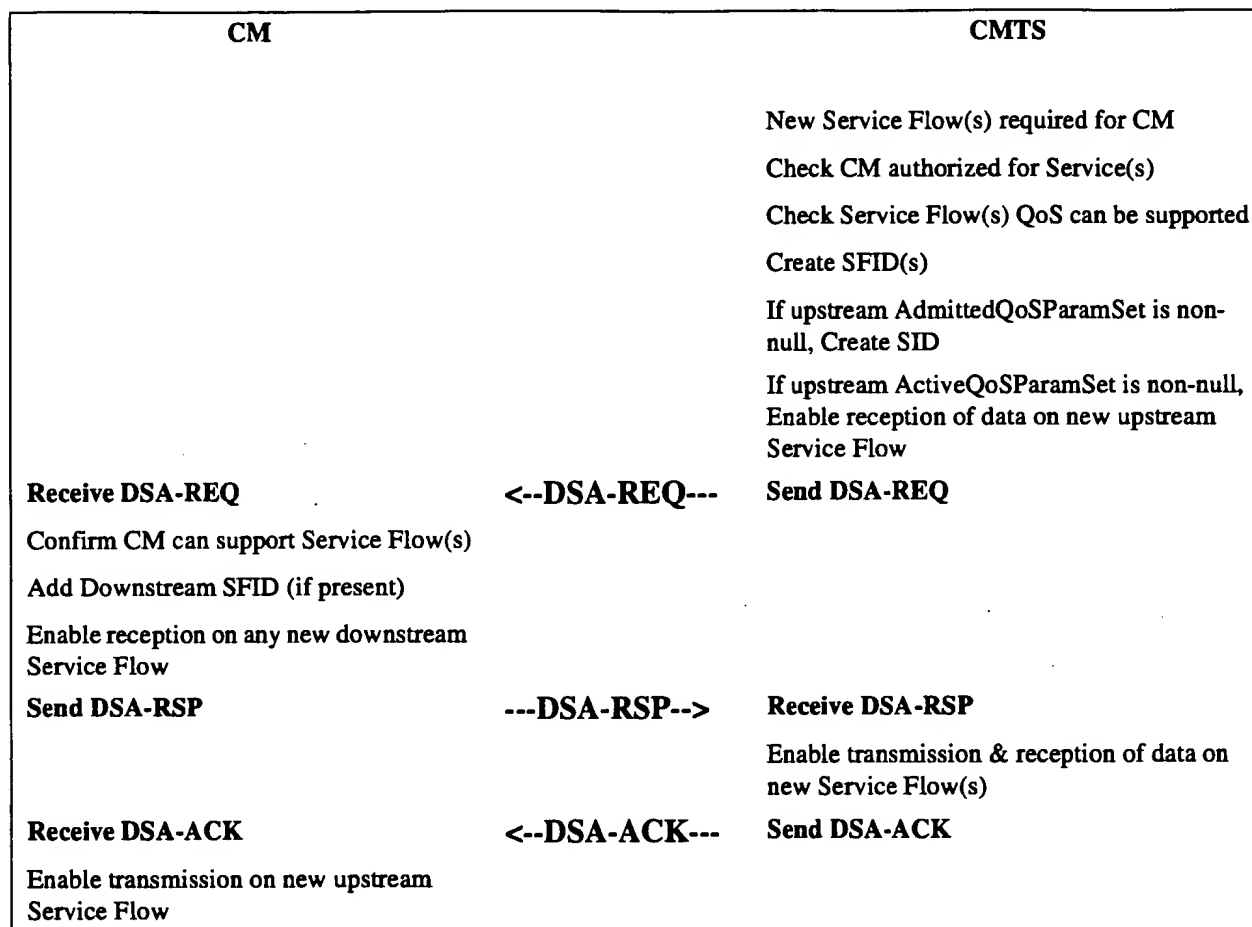


Figure 9-22. Dynamic Service Addition Initiated from CMTS (Figure edited per rfi-n-99048 06/30/99. ew)

1. Sentence edited per rfi-n-99048 06/30/99. ew

9.4.1.3 Dynamic Service Addition State Diagrams

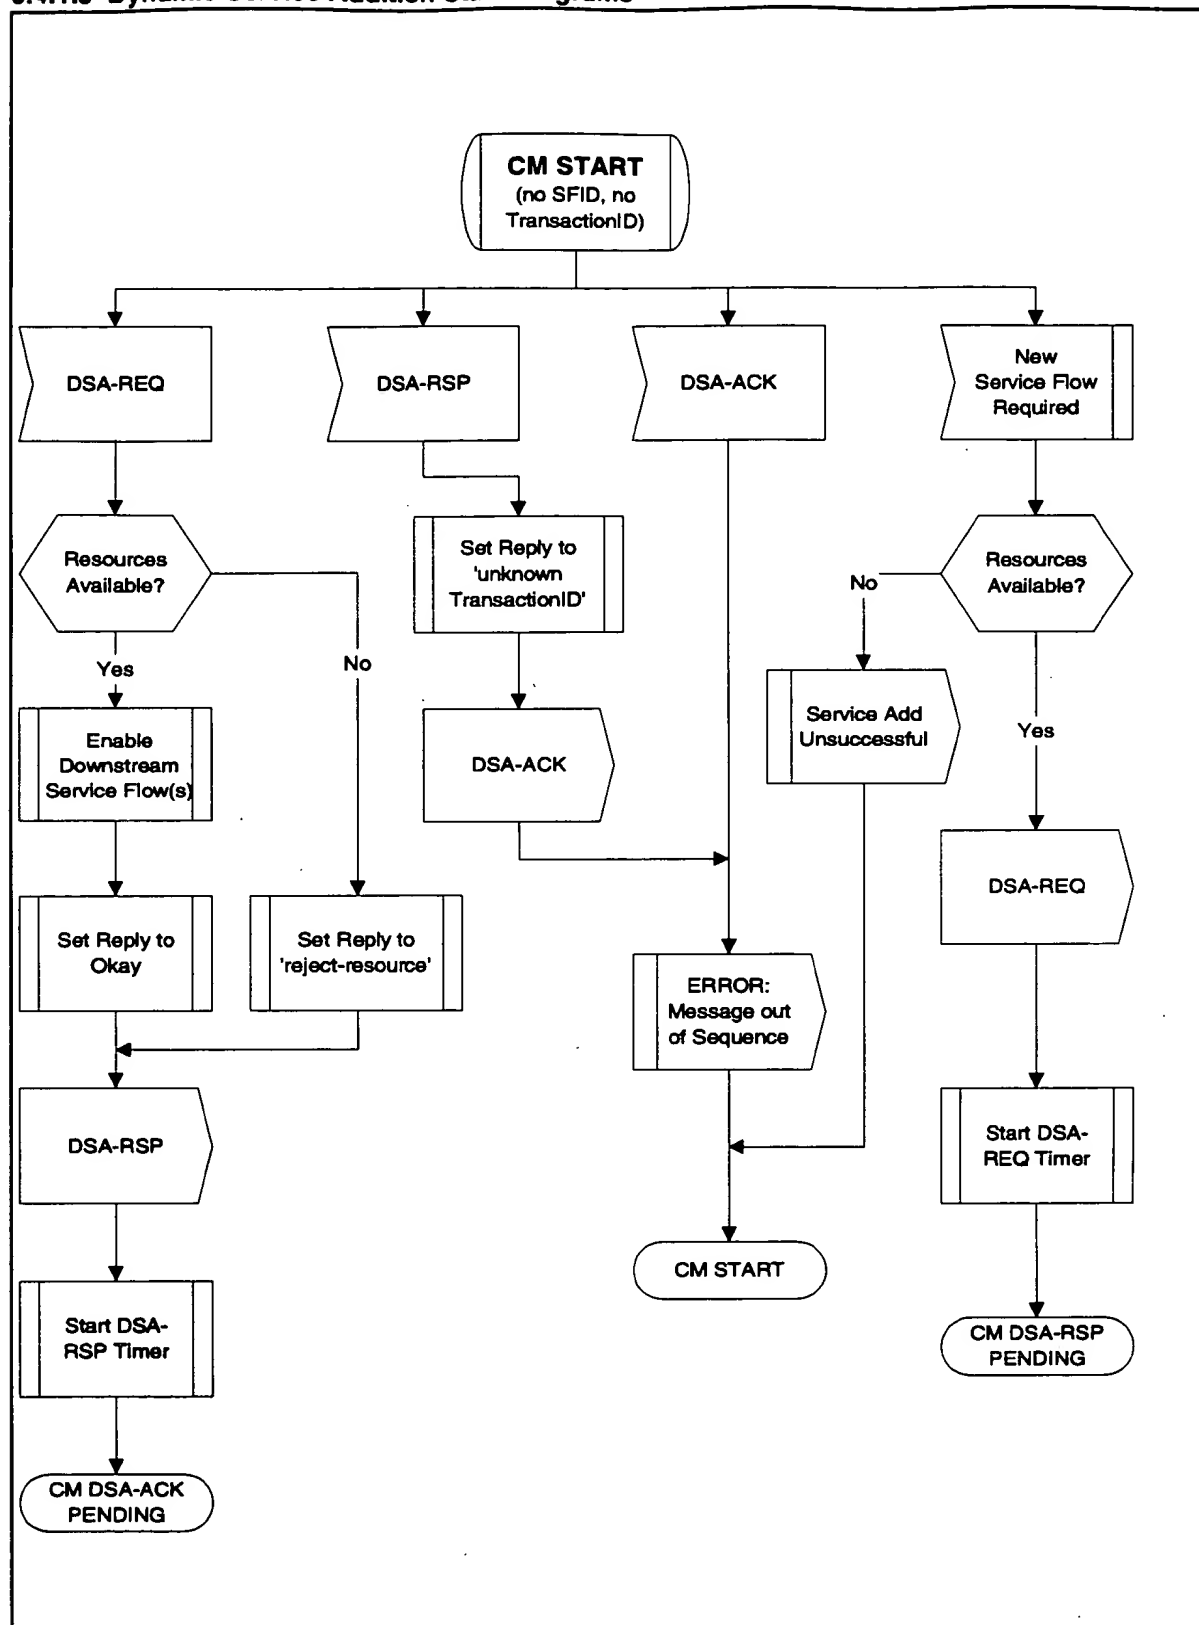


Figure 9-23. CM START State (DSA Transactions) (fig replaced 06/22/99 rfi-n-99043.ew)

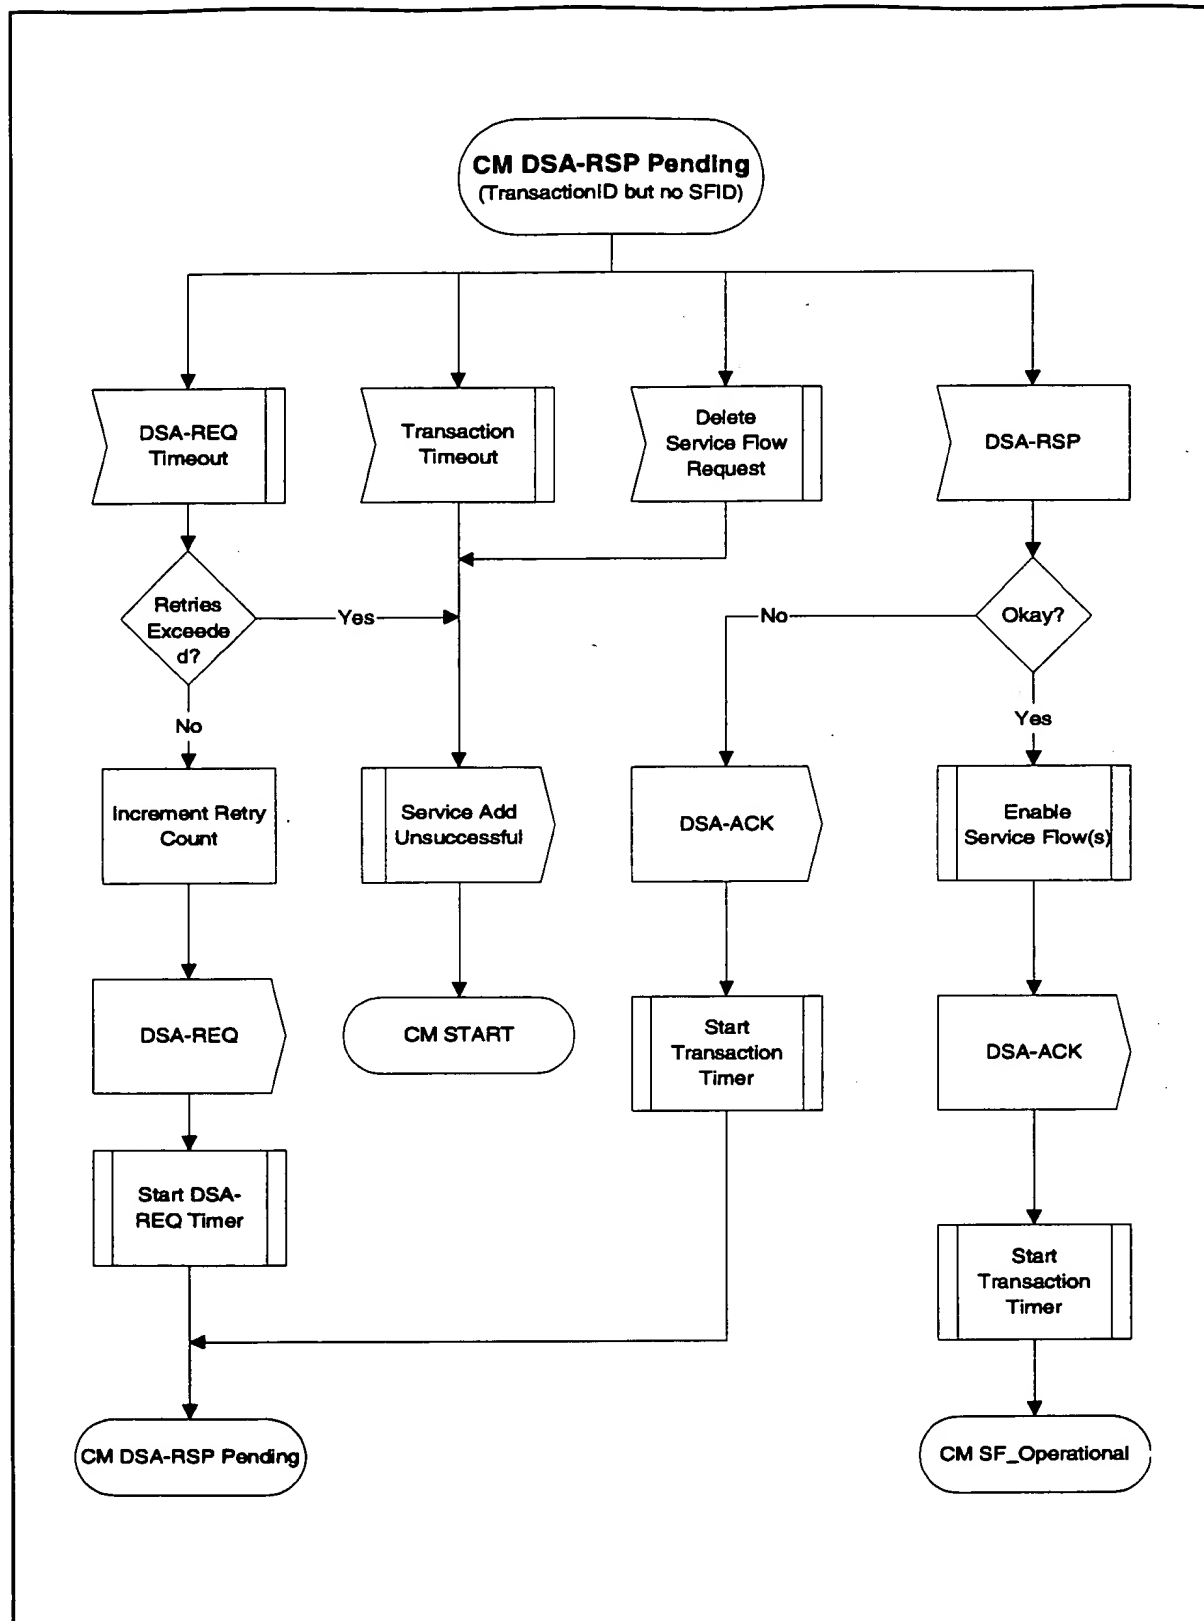


Figure 9-24. CM DSA-RSP Pending State (DSA Transactions)

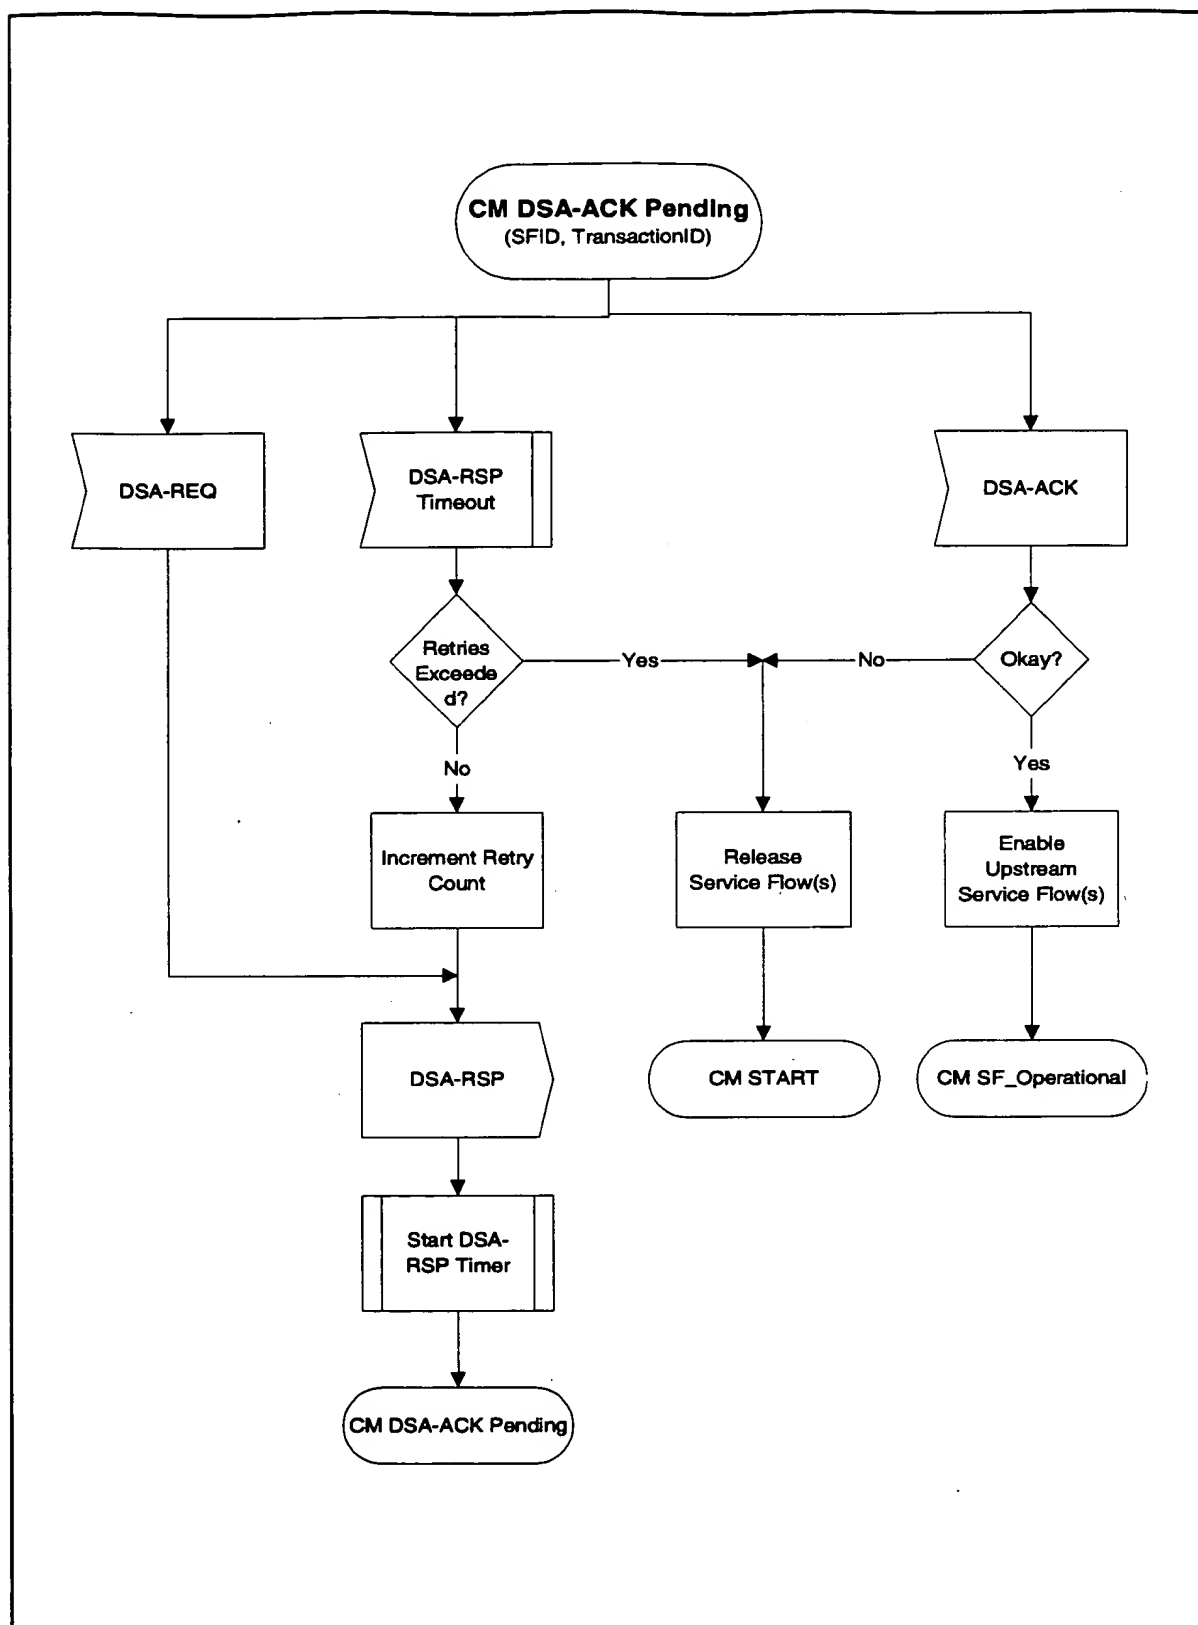


Figure 9-25. CM DSA-ACK Pending State (DSA Transactions)

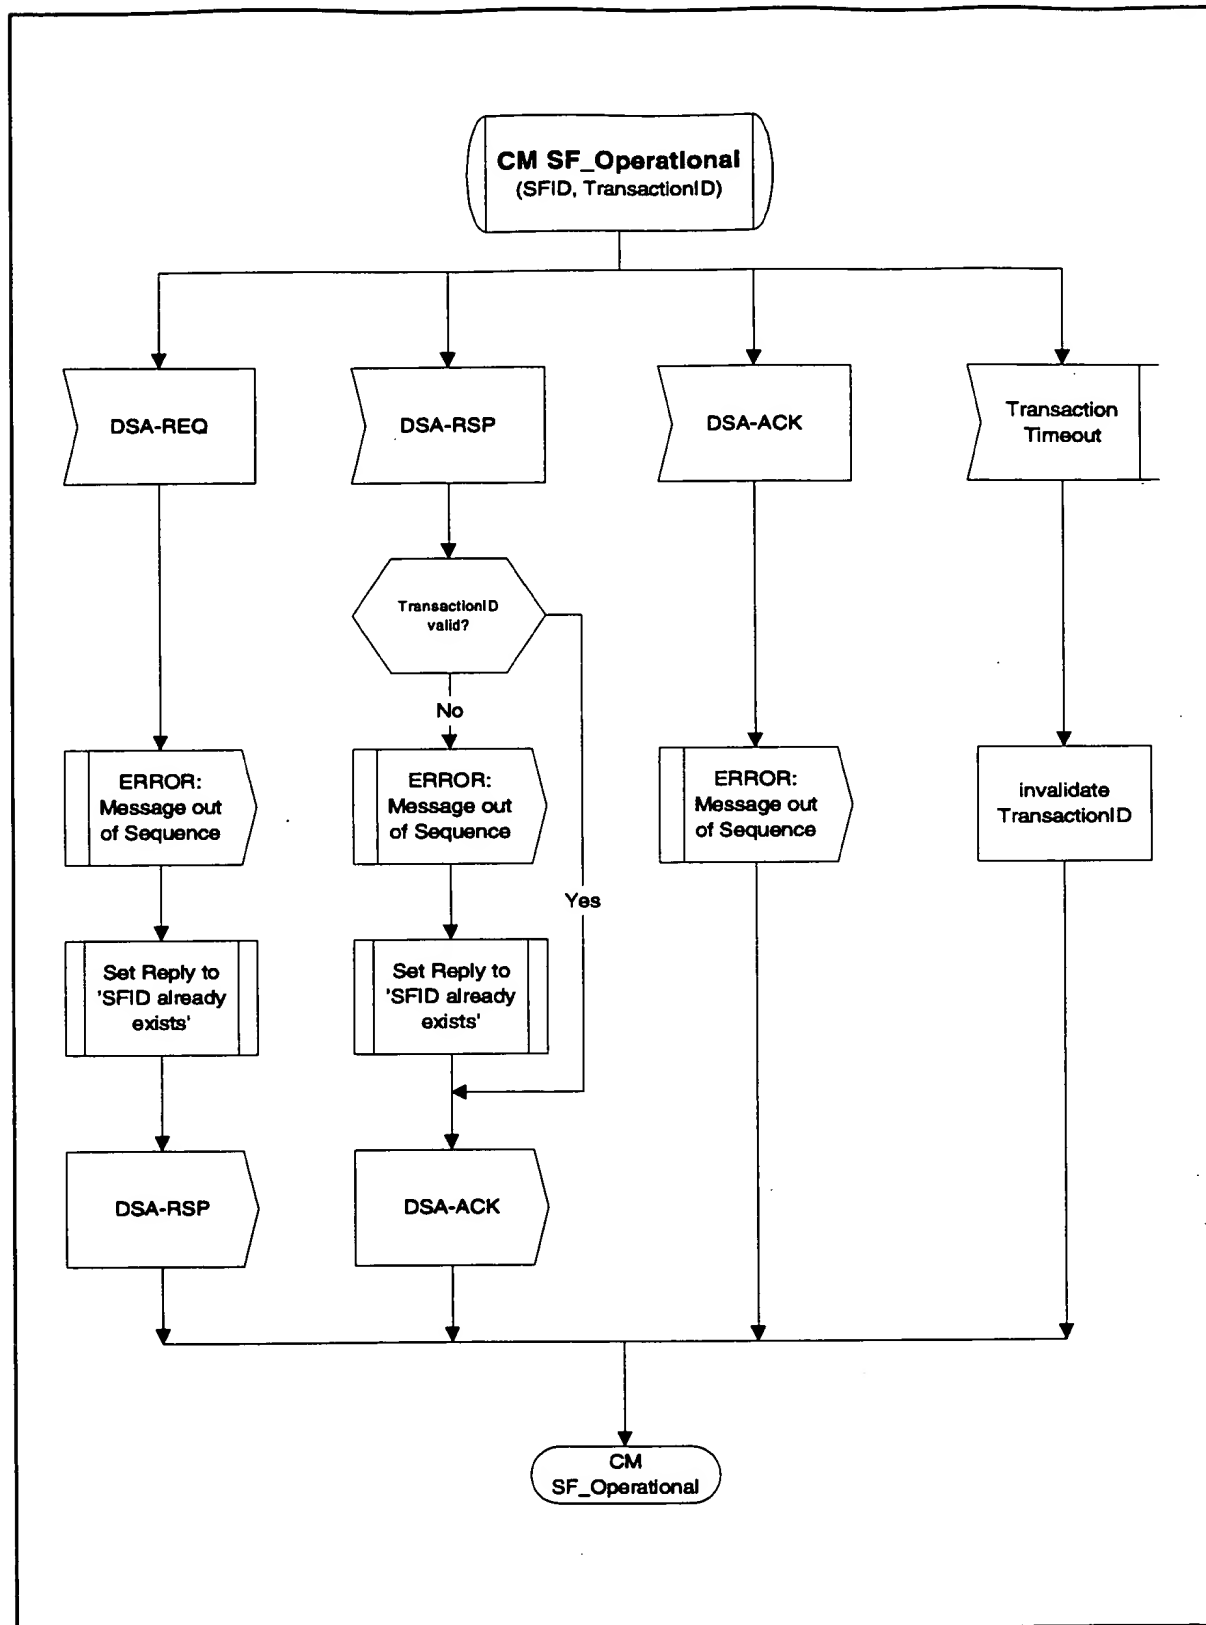


Figure 9-26. CM SF_Operational State (DSA Transactions)

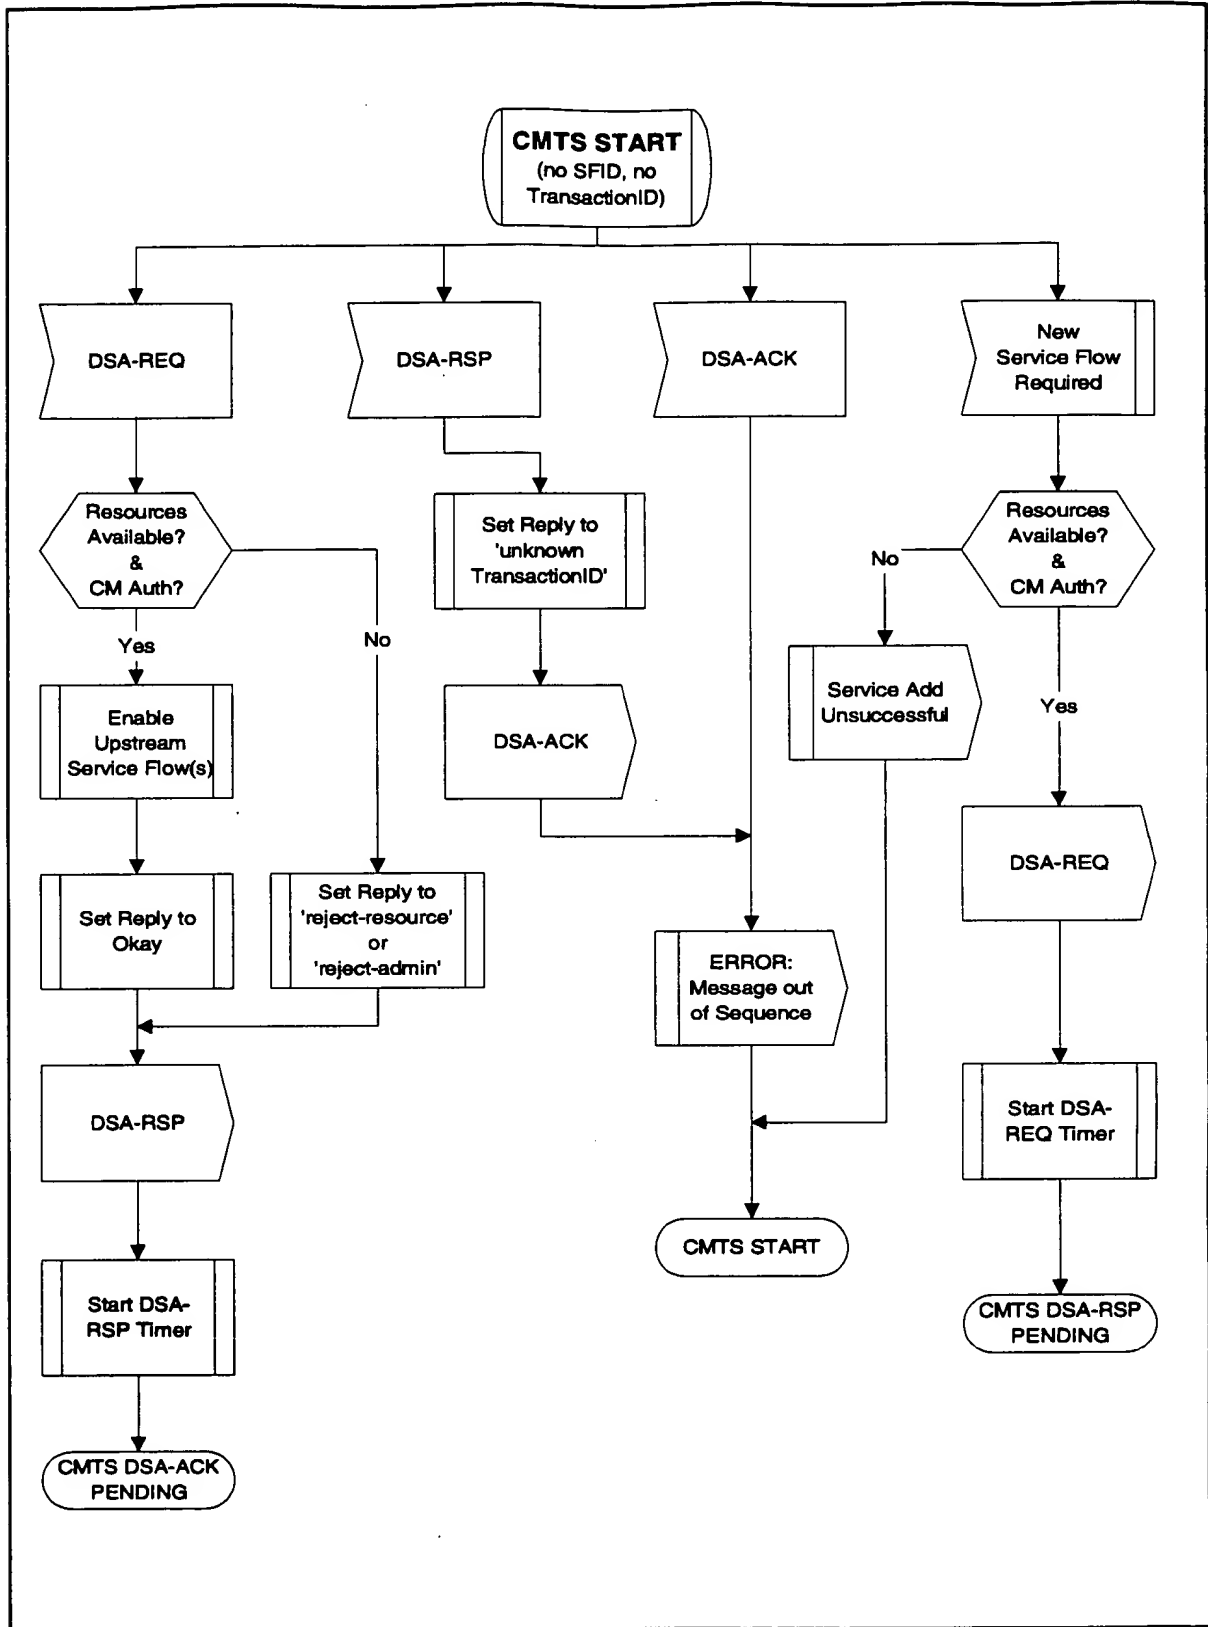


Figure 9-27. CMTS START State (DSA Transactions)

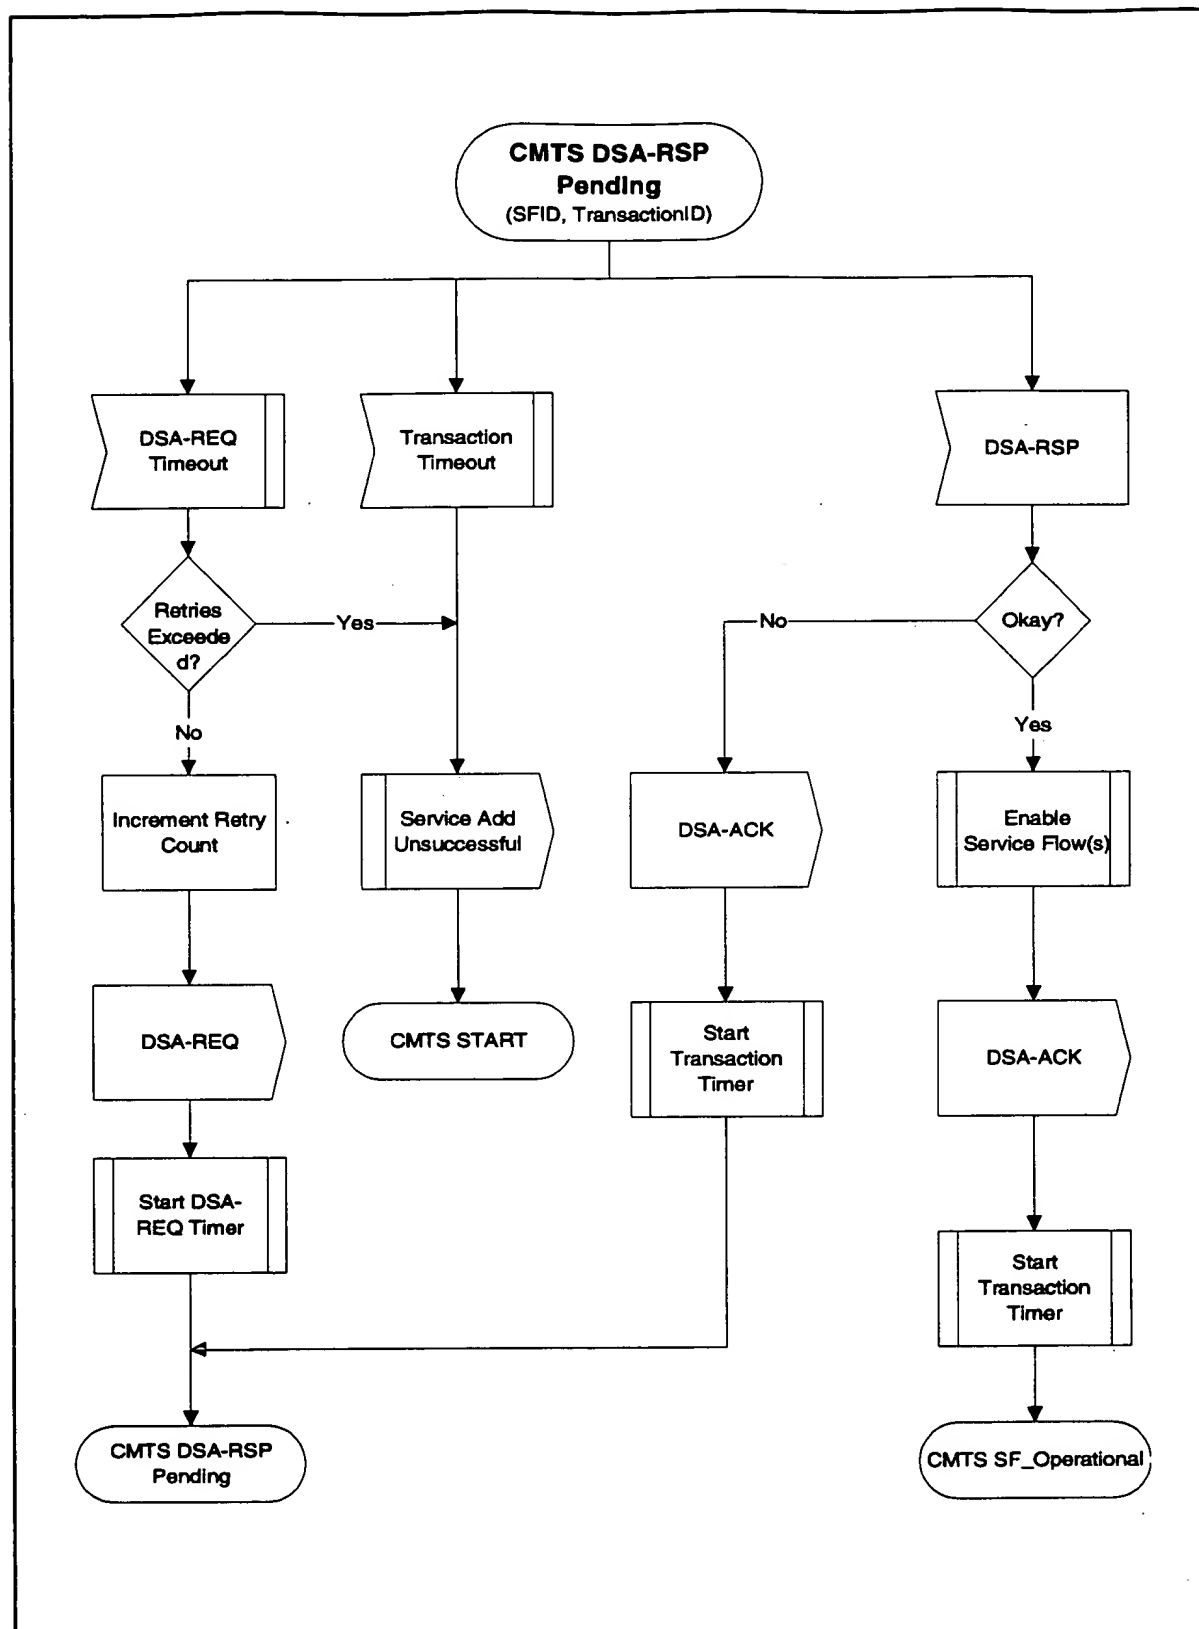


Figure 9-28. CMTS DSA-RSP Pending State (DSA Transactions)

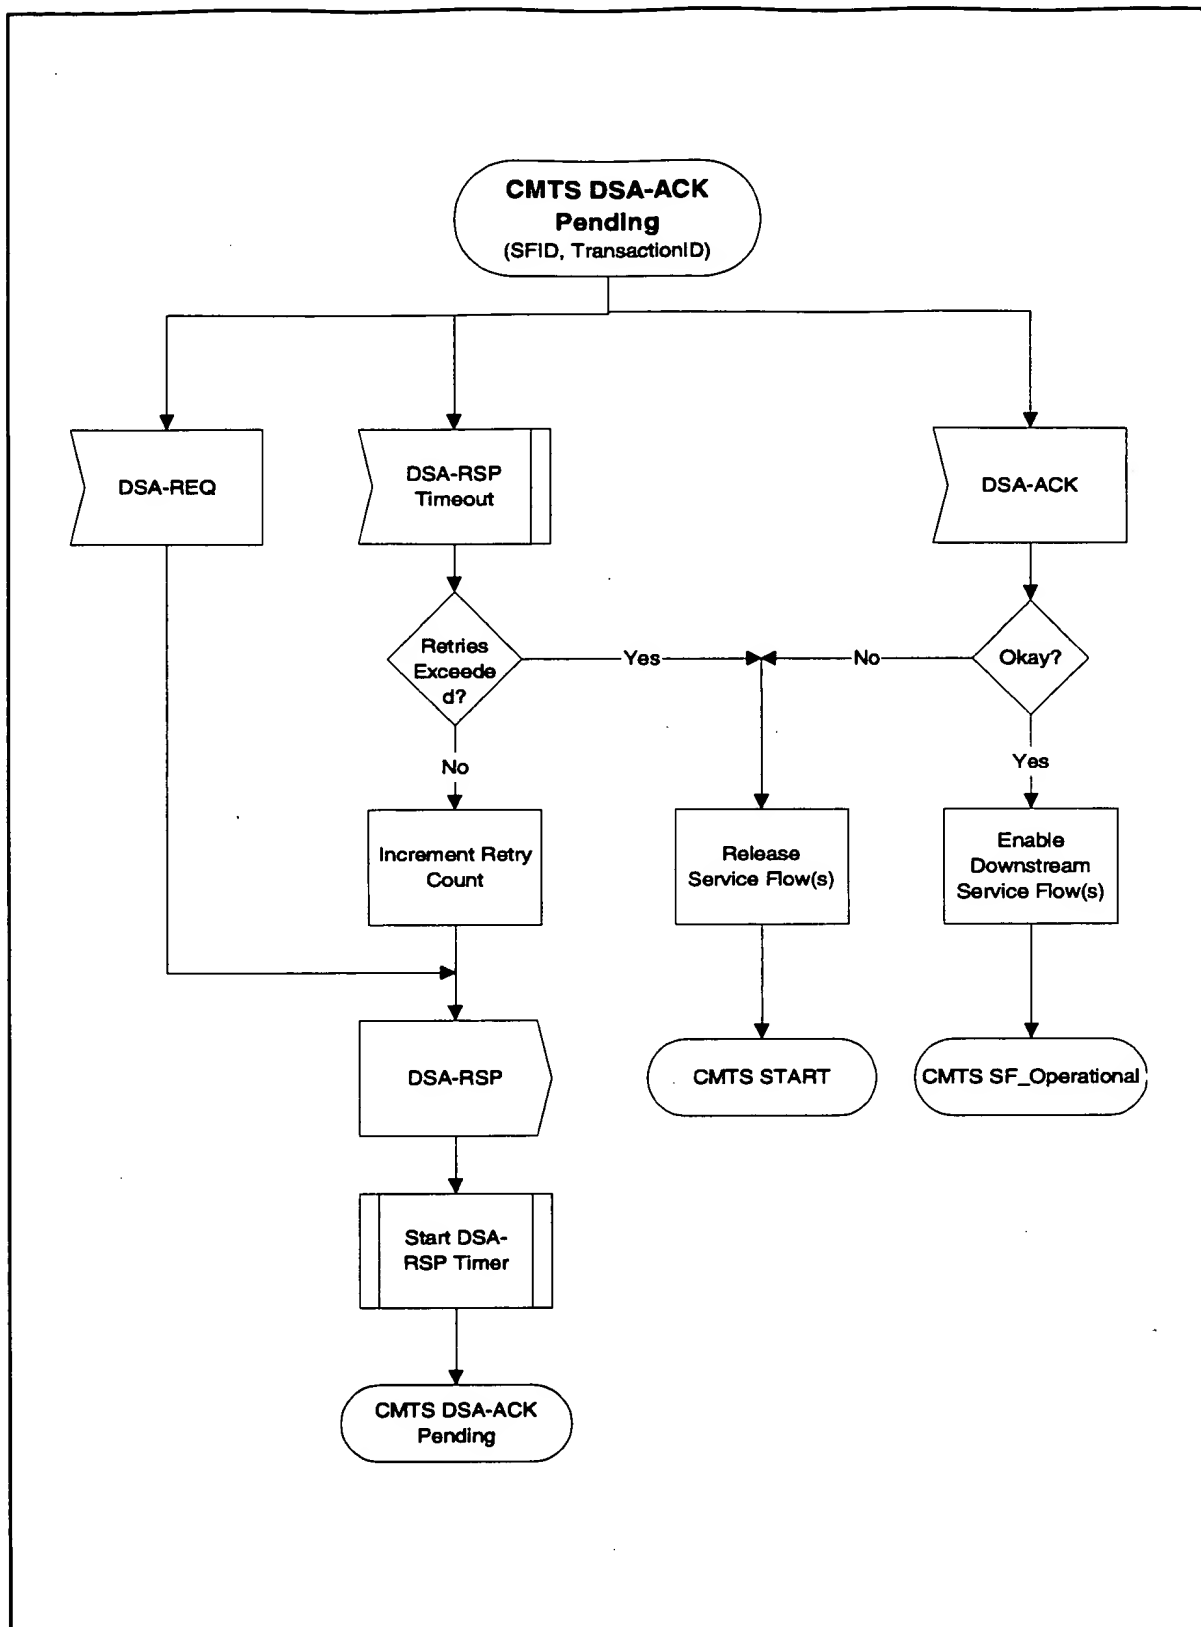


Figure 9-29. CMTS DSA-ACK Pending State (DSA Transactions)

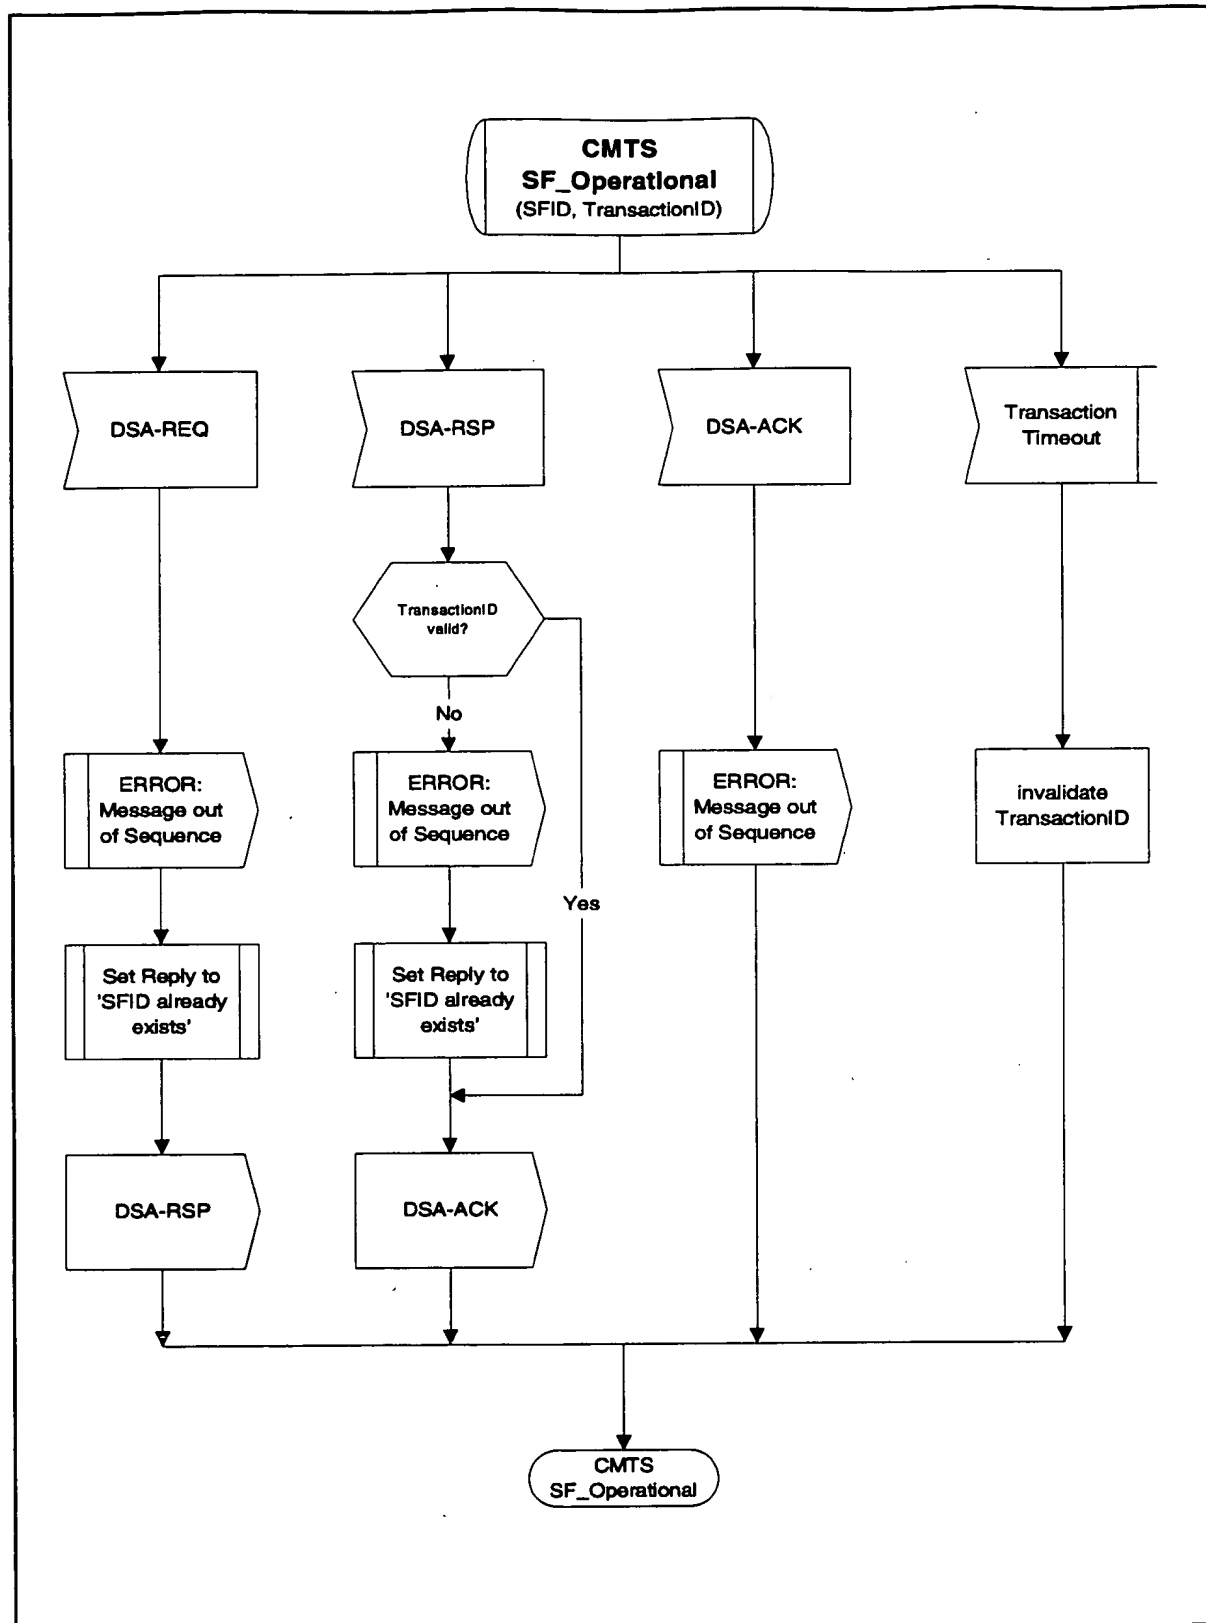


Figure 9-30. CMTS SF_Operational State (DSA Transactions)

9.4.2 Dynamic Service Change

The Dynamic Service Change (DSC) set of messages is used to modify the flow parameters associated with a Service Flow. Specifically, DSC can:

- Modify the Service Flow Specification
- Add, Delete or Replace a Flow Classifier
- Add, Delete or Set PHS elements¹

A single DSC message exchange can modify the parameters of one downstream service flow and/or one upstream service flow.²

To prevent packet loss, any required bandwidth changes are sequenced between the CM and the CMTS. As an example, in the upstream, if the Service Flow bandwidth is to be reduced, the CM reduces its payload bandwidth first, and then the CMTS reduces the bandwidth scheduled for the Service Flow³. If the Service Flow bandwidth is to be increased, the CMTS increases the bandwidth scheduled for the Service Flow first, and then the CM increases its payload bandwidth.

If the bandwidth changes are complex, it may not be obvious to the CM when to effect the bandwidth changes. This information may be signalled to the CM from a higher layer entity. Similarly, if the DSC signaling is initiated by the CMTS, the CMTS MAY indicate to the CM whether it should install or remove Classifiers upon receiving the DSC-Request or whether it should postpone this installation until receiving the DSC-Ack (refer to C.2.1.8)

Any service flow can be deactivated with a Dynamic Service Change command by sending a DSC-REQ message, referencing the Service Flow Identifier, and including a null ActiveQoSParameterSet⁴. However, if a Primary Service Flow of a CM is deactivated that CM is de-registered and MUST re-register. Therefore, care should be taken before deactivating such Service Flows. If a Service Flow that was provisioned during registration is deactivated, the provisioning information for that Service Flow MUST be maintained until the Service Flow is reactivated.

A CM MUST have only one DSC transaction outstanding per Service Flow. If it detects a second transaction initiated by the CMTS, the CM MUST abort the transaction it initiated and allow the CMTS initiated transaction to complete.

A CMTS MUST have only one DSC transaction outstanding per Service Flow. If it detects a second transaction initiated by the CM, the CMTS MUST abort the transaction the CM initiated and allow the CMTS initiated transaction to complete.

Note: Currently anticipated applications would probably control a Service Flow through either the CM or CMTS, and not both. Therefore the case of a DSC being initiated simultaneously by the CM and CMTS is considered as an exception condition and treated as one.

1. Bullet edited 06/22/99 per rfi-n-99043. ew

2. Paragraph added 06/30/99 per rfi-n-99048. ew

3. Sentence edited 06/22/99 per rfi-n-99043. ew

4. Sentence edited 06/22/99 per rfi-n-99043. ew

9.4.2.1 CM-Initiated Dynamic Service Change

A CM that needs to change a Service Flow definition performs the following operations.

The CM informs the CMTS using a Dynamic Service Change Request message (DSC-REQ). The CMTS **MUST** decide if the referenced Service Flow can support this modification. The CMTS **MUST** respond with a Dynamic Service Change Response (DSC-RSP) indicating acceptance or rejection. The CM reconfigures the Service Flow if appropriate, and then **MUST** respond with a Dynamic Service Change Acknowledge (DSC-ACK).

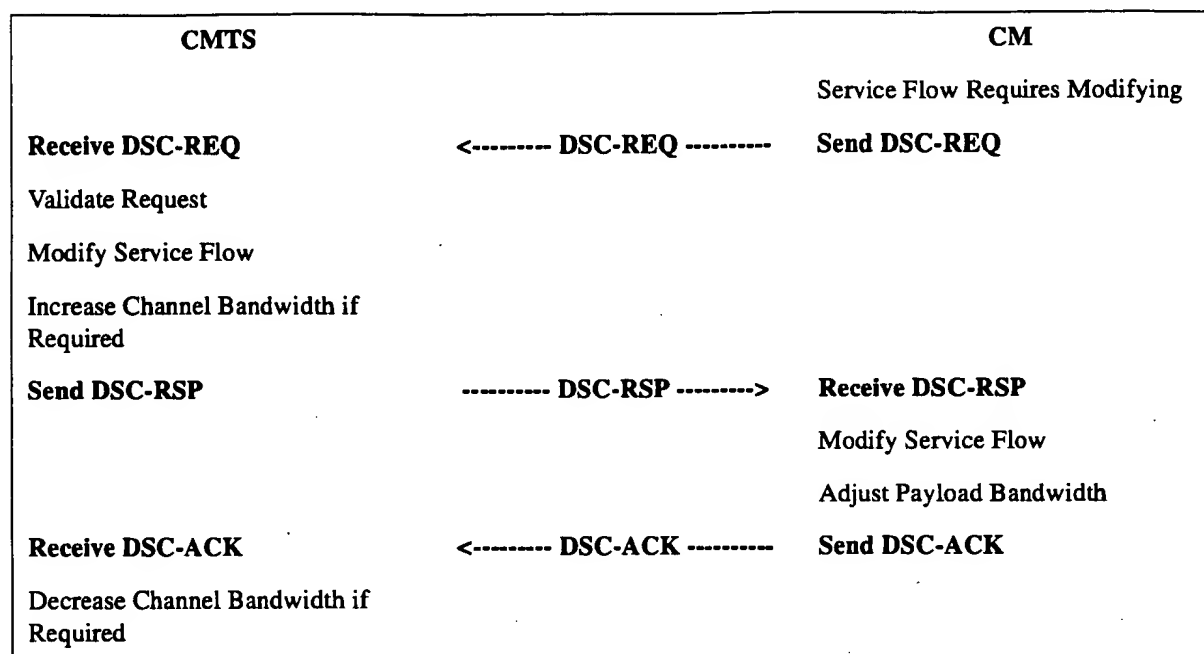


Figure 9-31. CM-Initiated DSC

9.4.2.2 CMTS-Initiated Dynamic Service Change

A CMTS that needs to change a Service Flow definition performs the following operations.

The CMTS **MUST** decide if the referenced Service Flow can support this modification. If so, the CMTS informs the CM using a Dynamic Service Change Request message (DSC-REQ). The CM checks that it can support the service change, and **MUST** respond using a Dynamic Service Change Response (DSC-RSP) indicating acceptance or rejection. The CMTS reconfigures the Service Flow if appropriate, and then **MUST** respond with a Dynamic Service Change Acknowledgment (DSC-ACK)

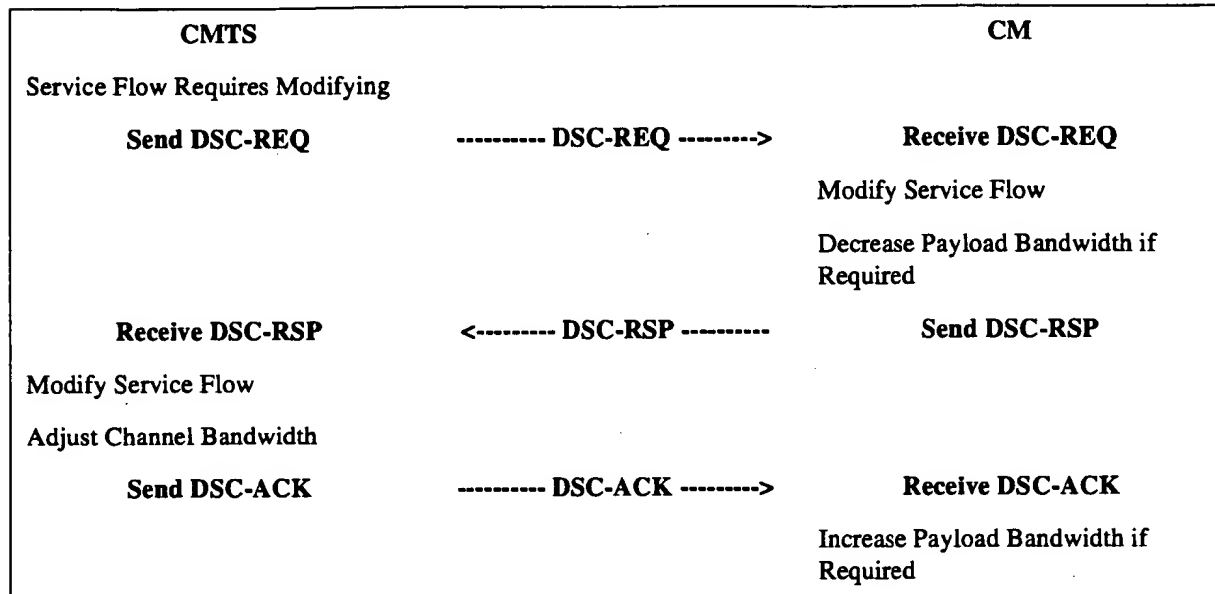


Figure 9-32. CMTS-Initiated DSC

9.4.2.3 Dynamic Service Change State Diagrams

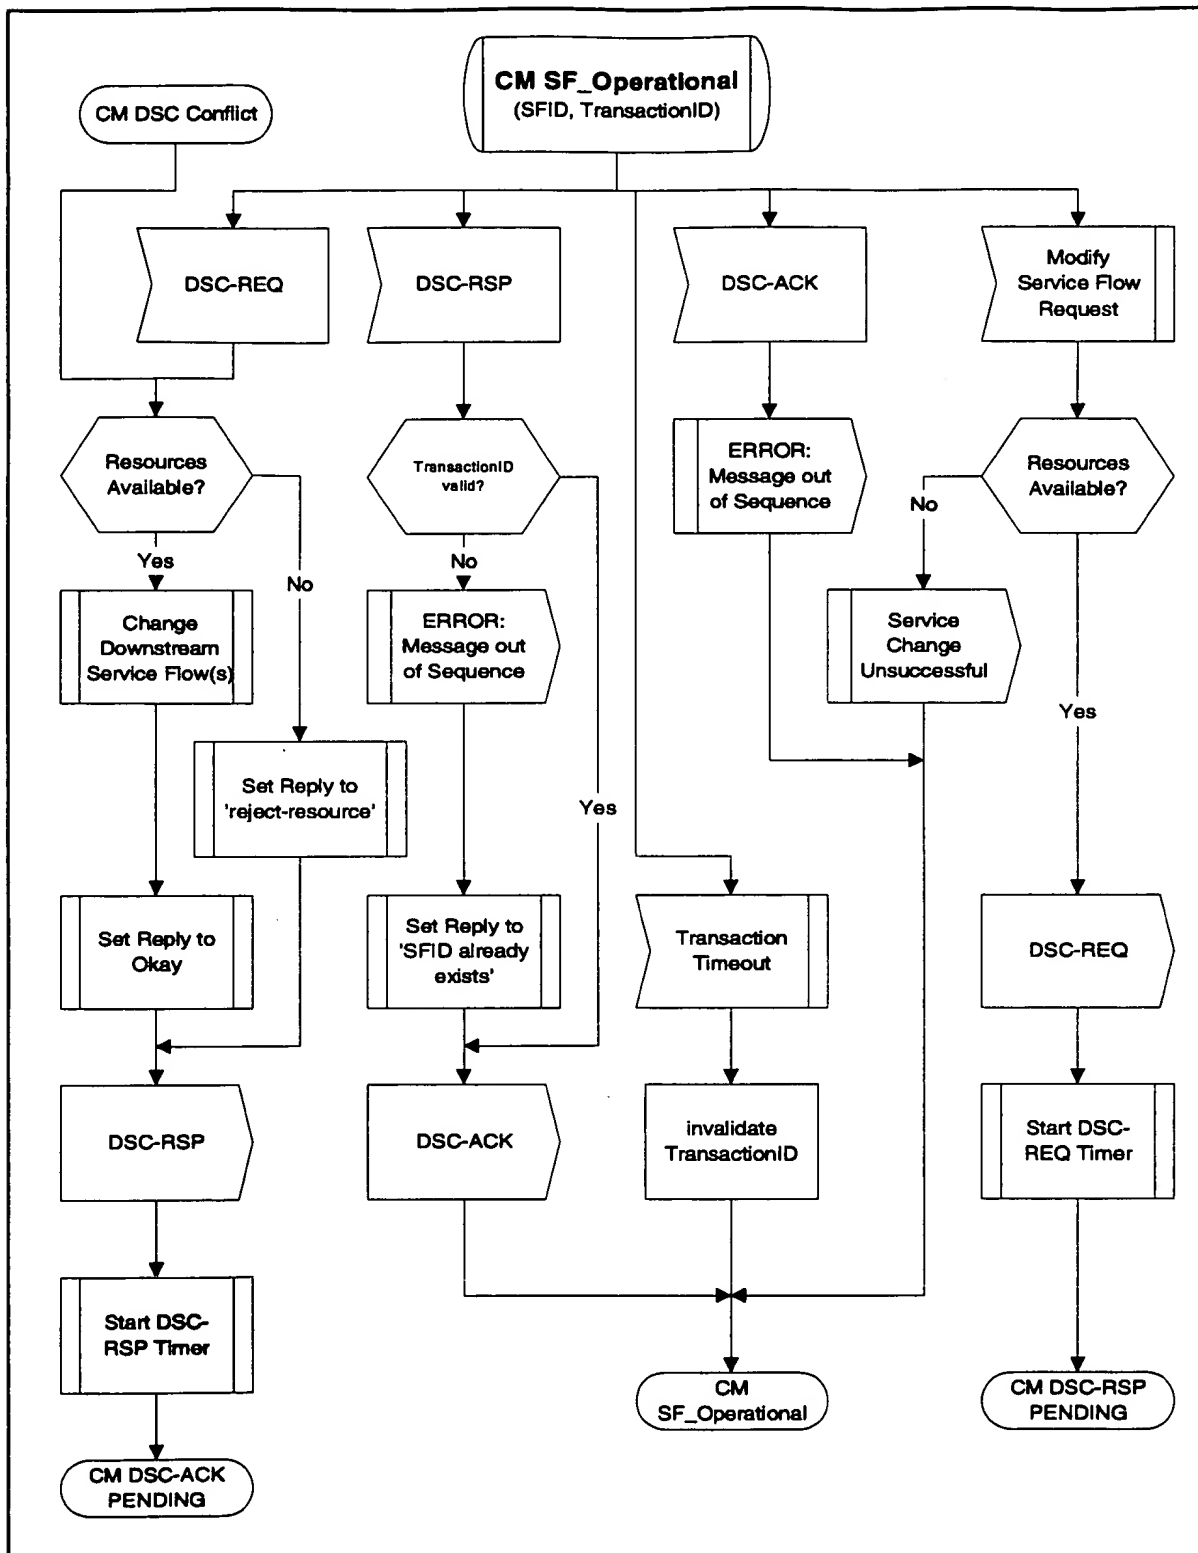


Figure 9-33. CM SF_Operational State (DSC Transactions)

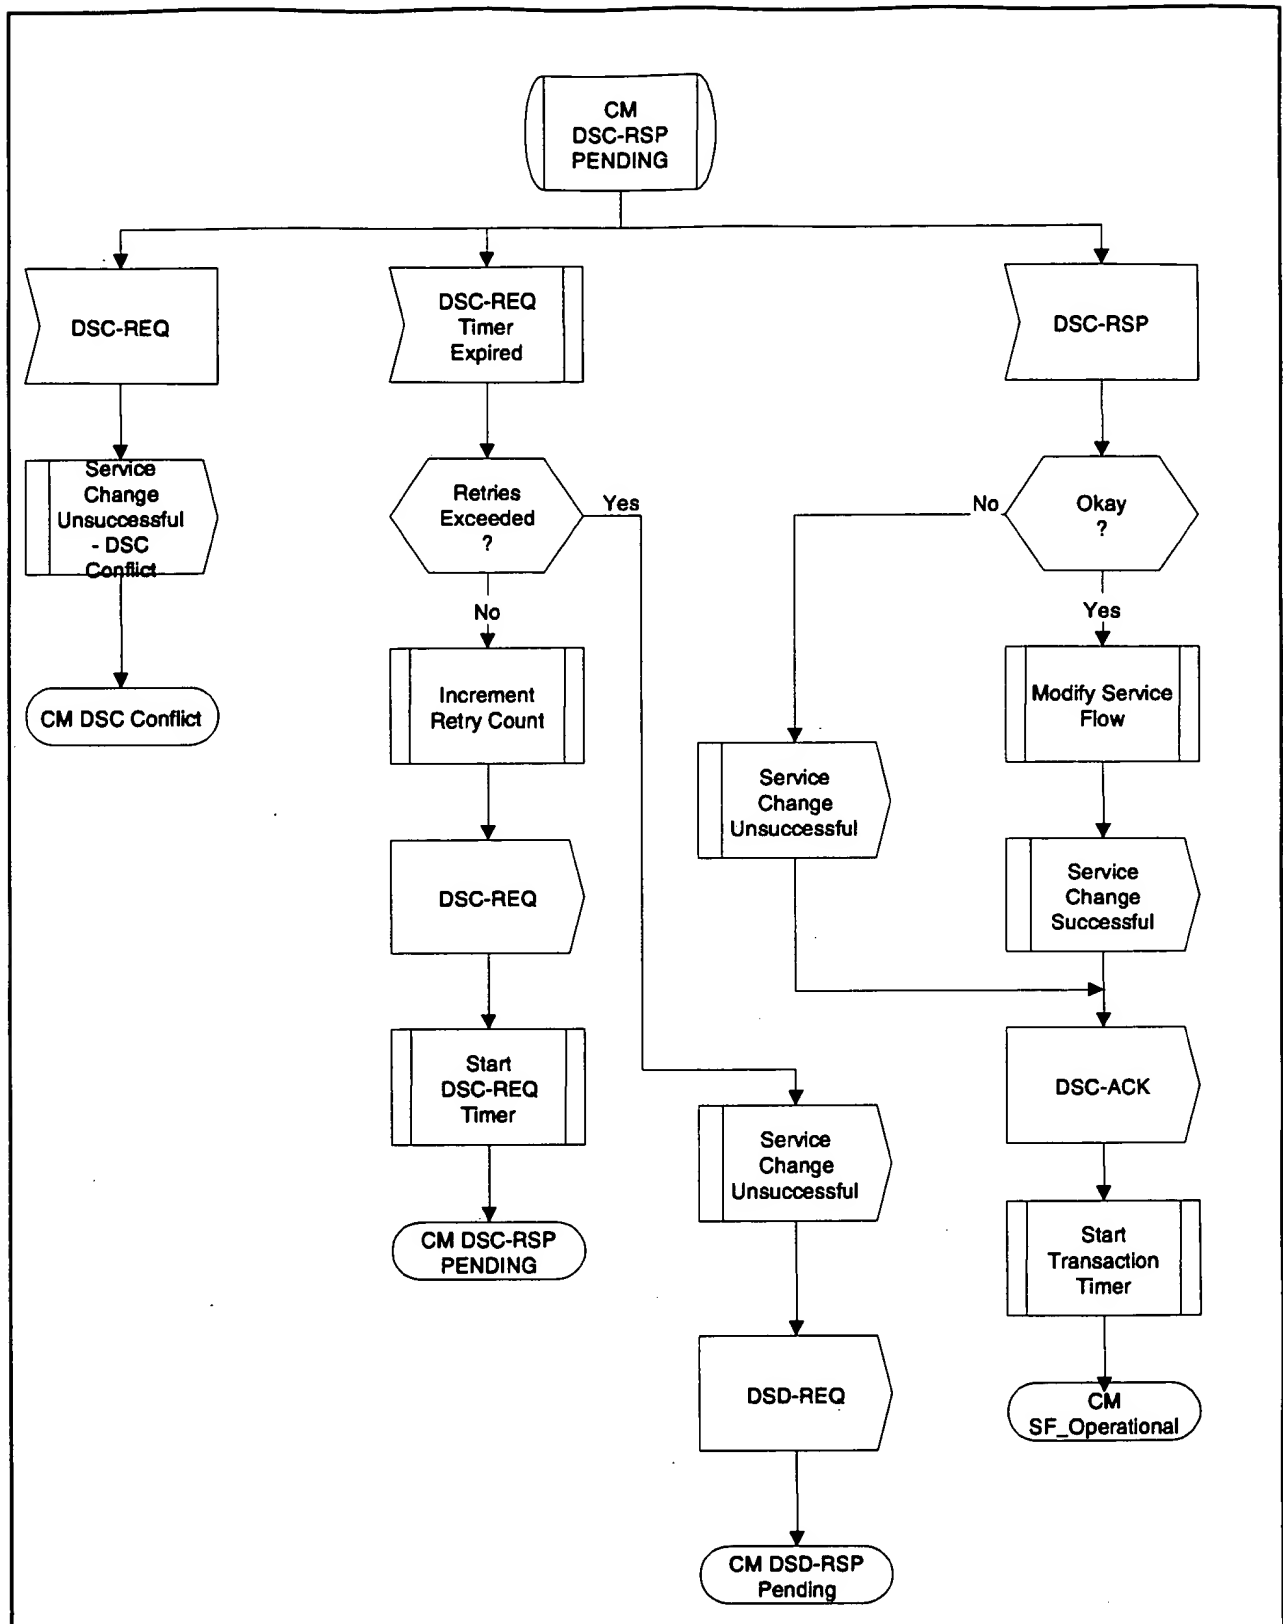


Figure 9-34. CM DSC-RSP Pending State (DSC Transactions) (figure edited 06/22/99 per rfi-n-99056.ew)

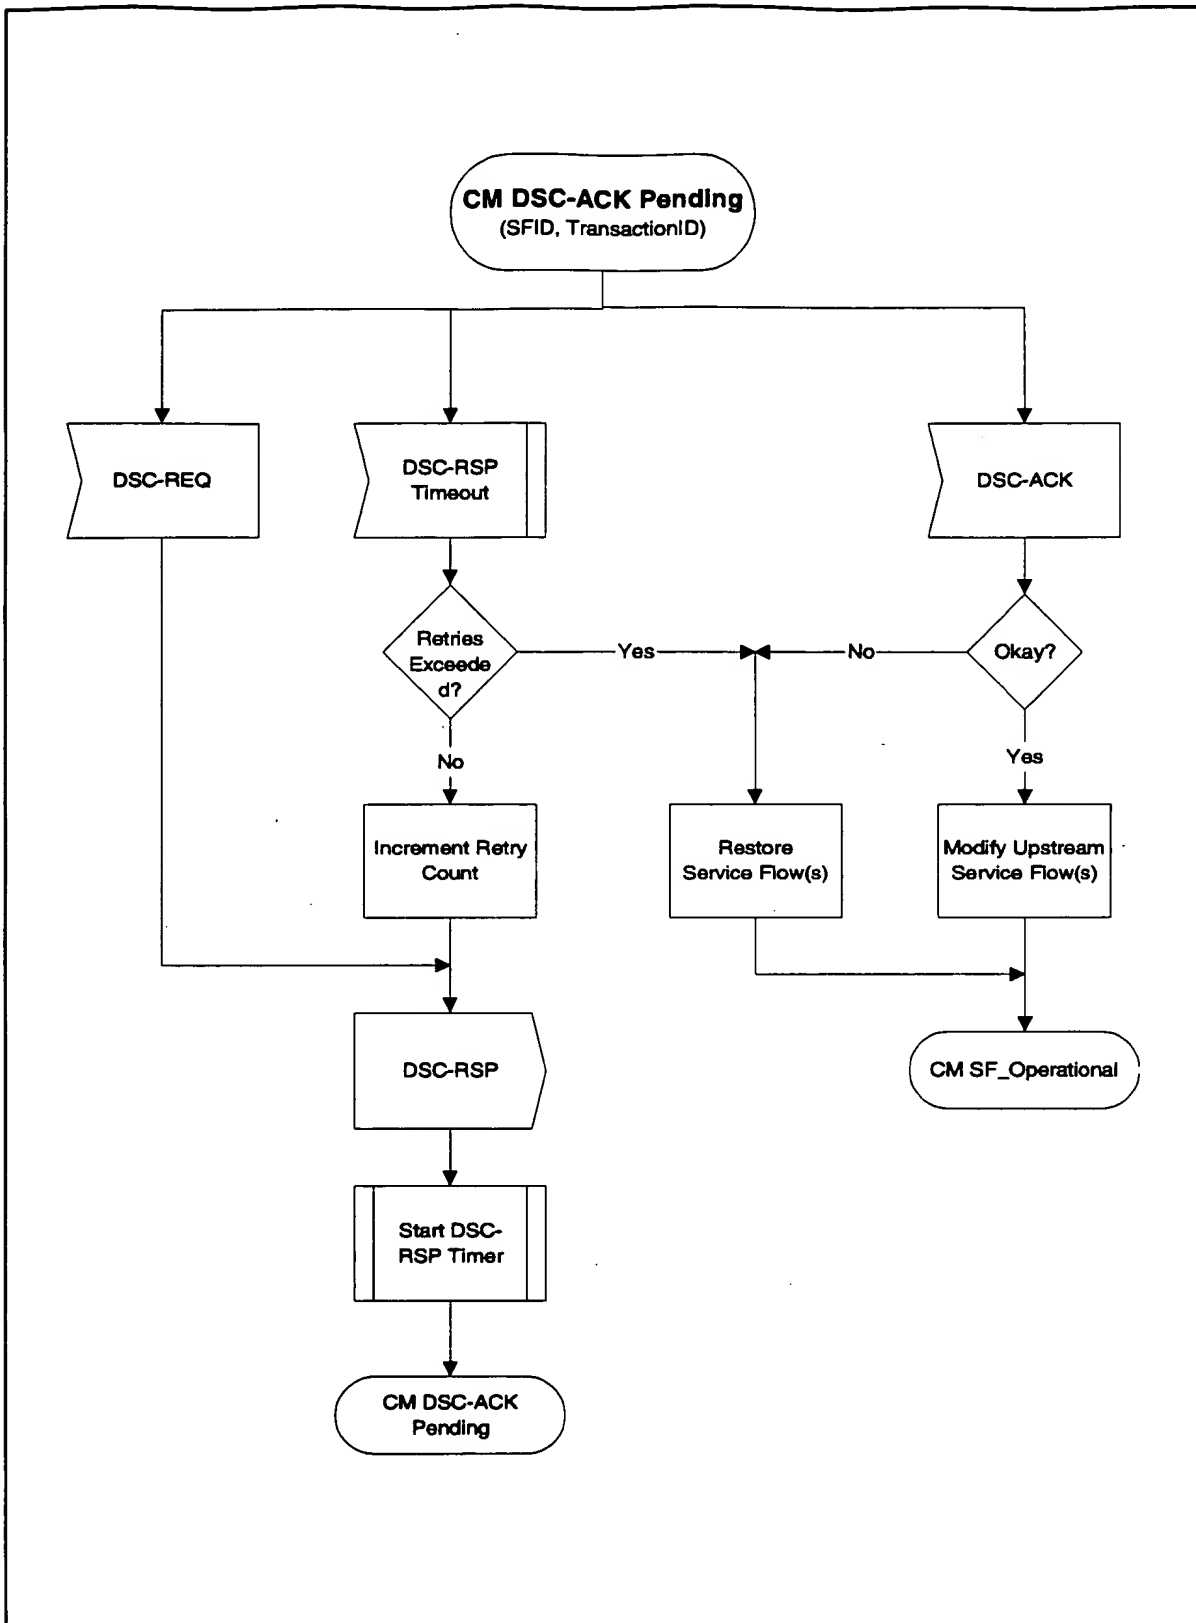


Figure 9-35. CM DSC-ACK Pending State (DSC Transactions)

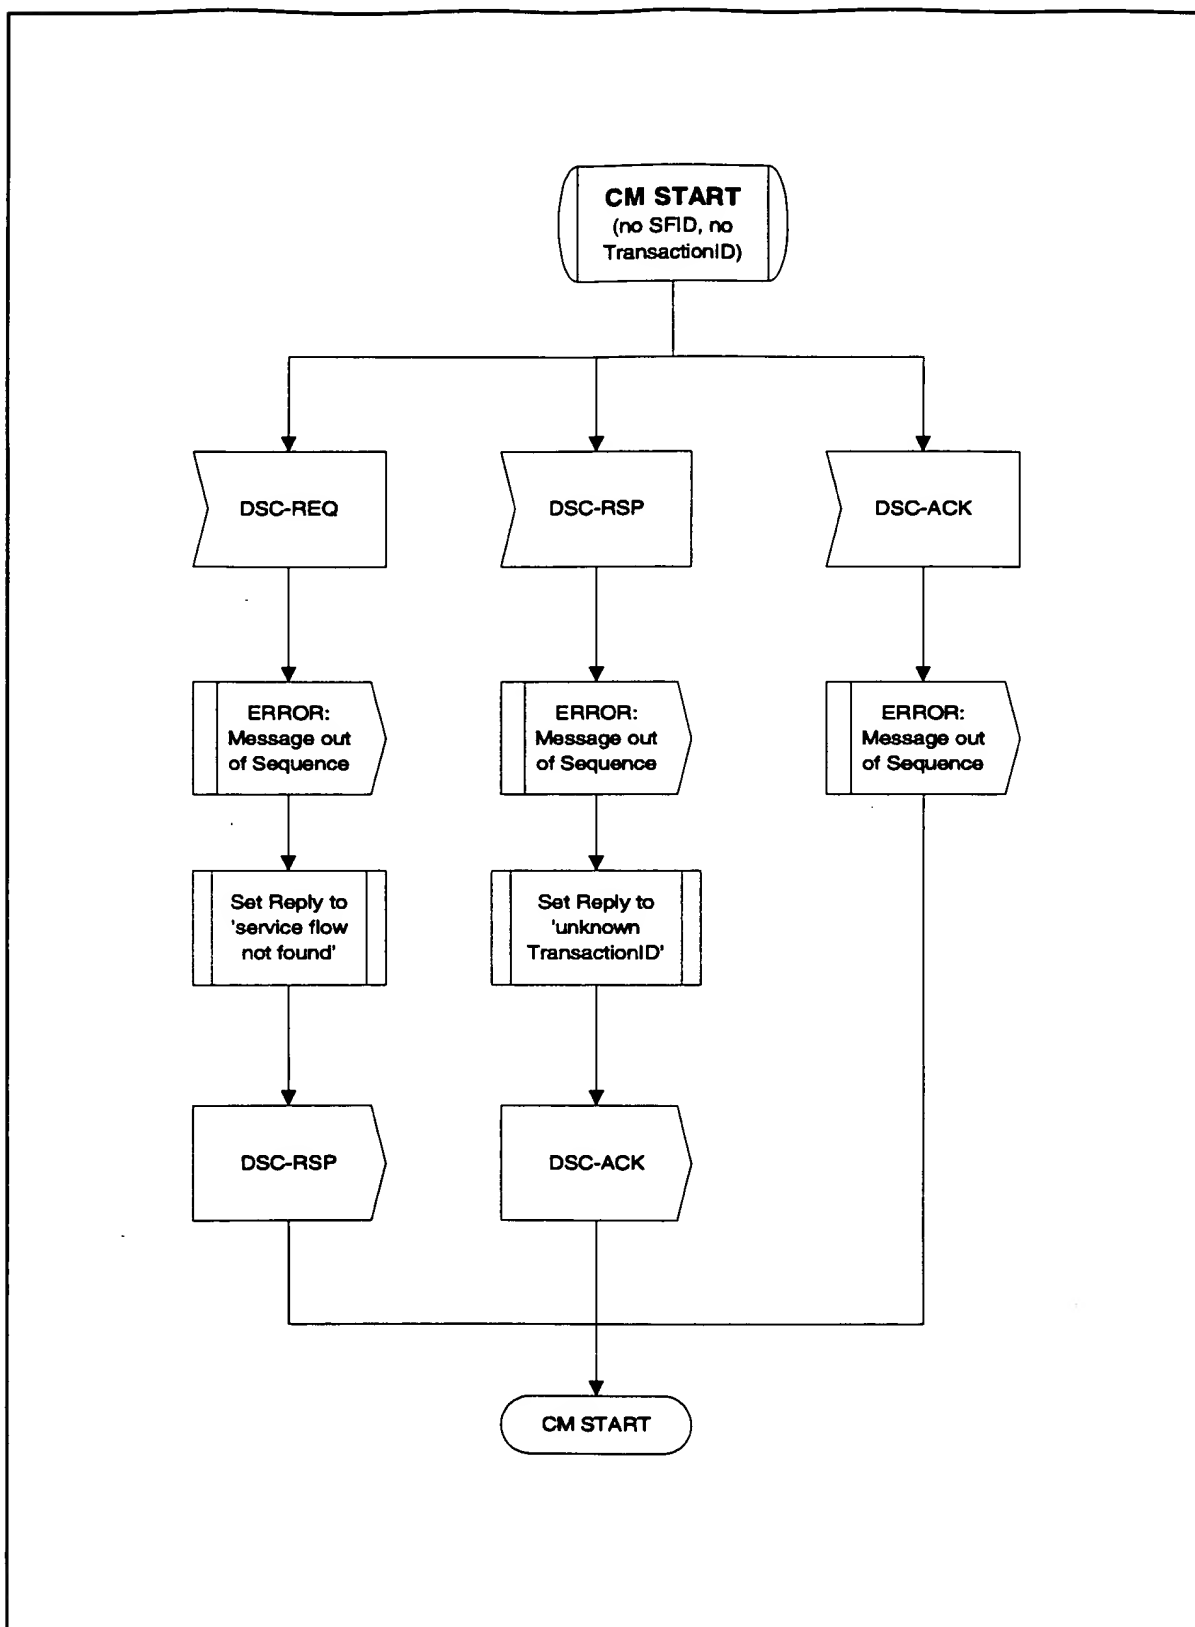


Figure 9-36. CM START State (DSC Transactions)

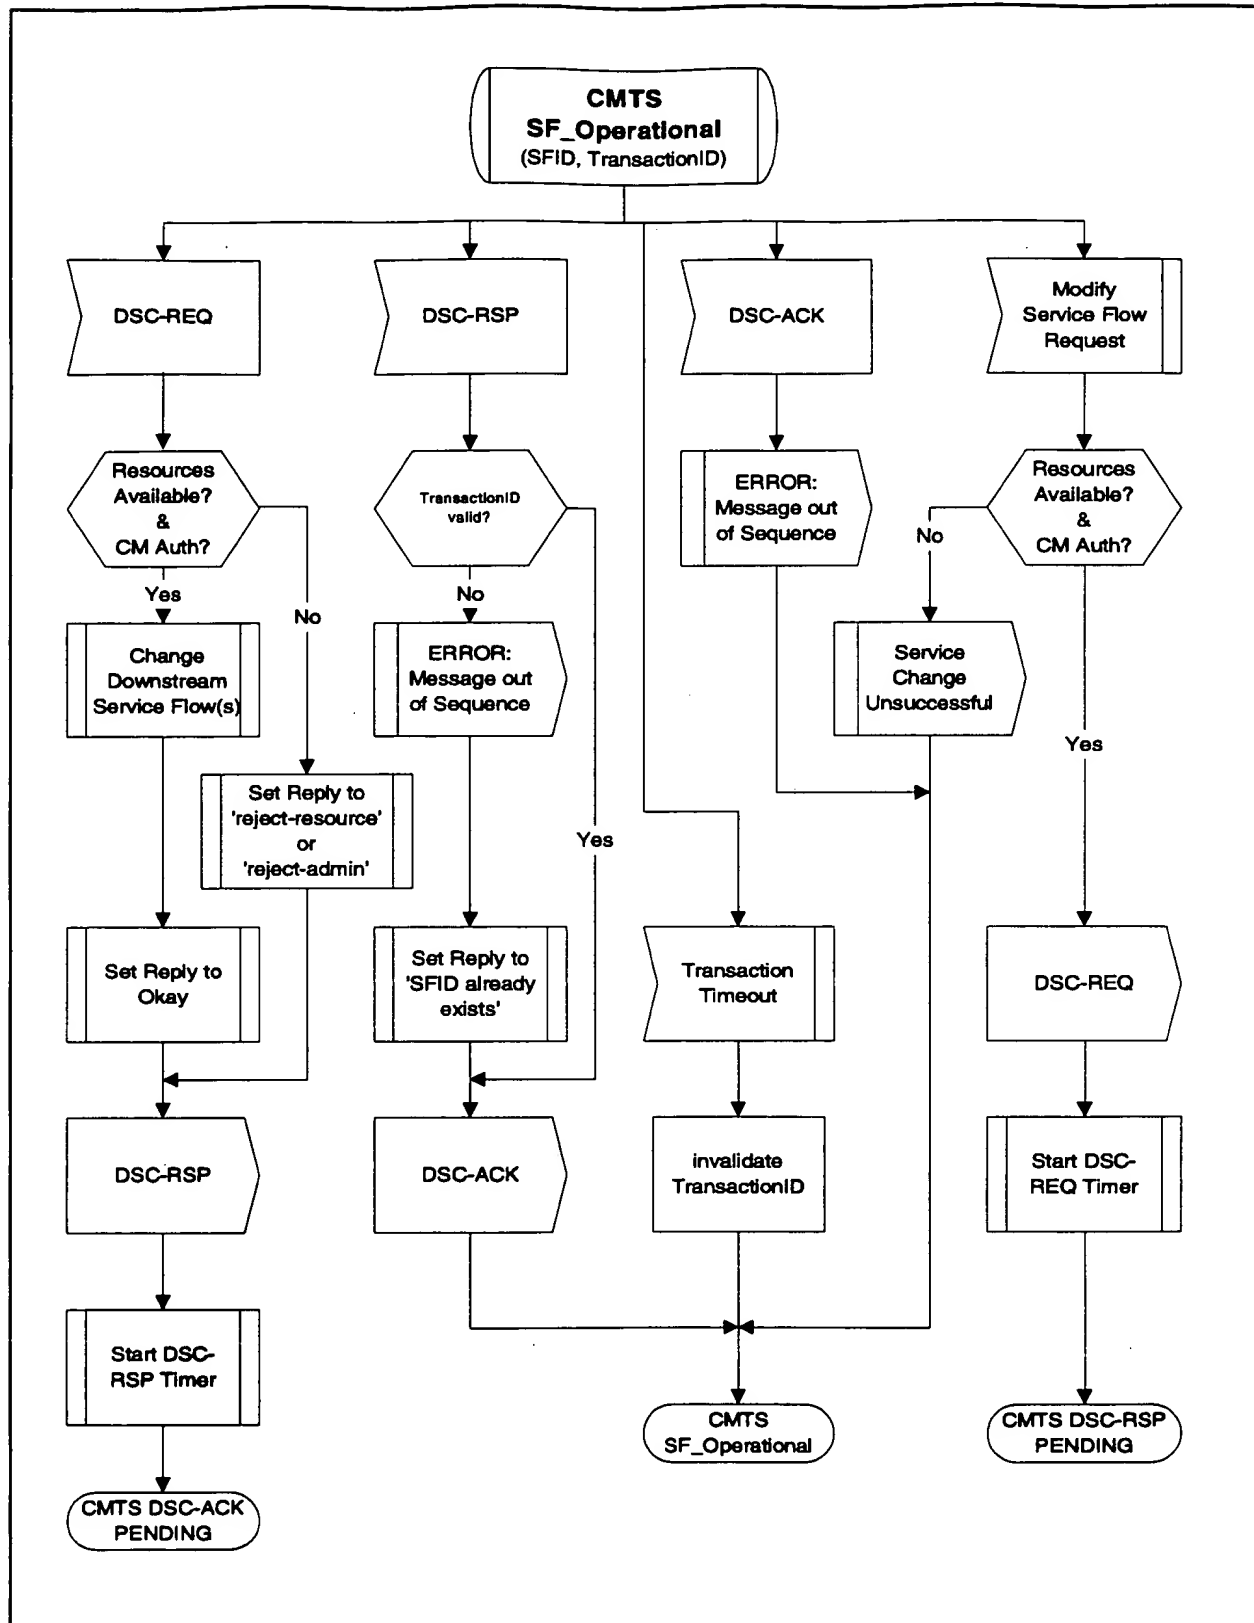


Figure 9-37. CMTS SF_Operational State (DSC Transactions)

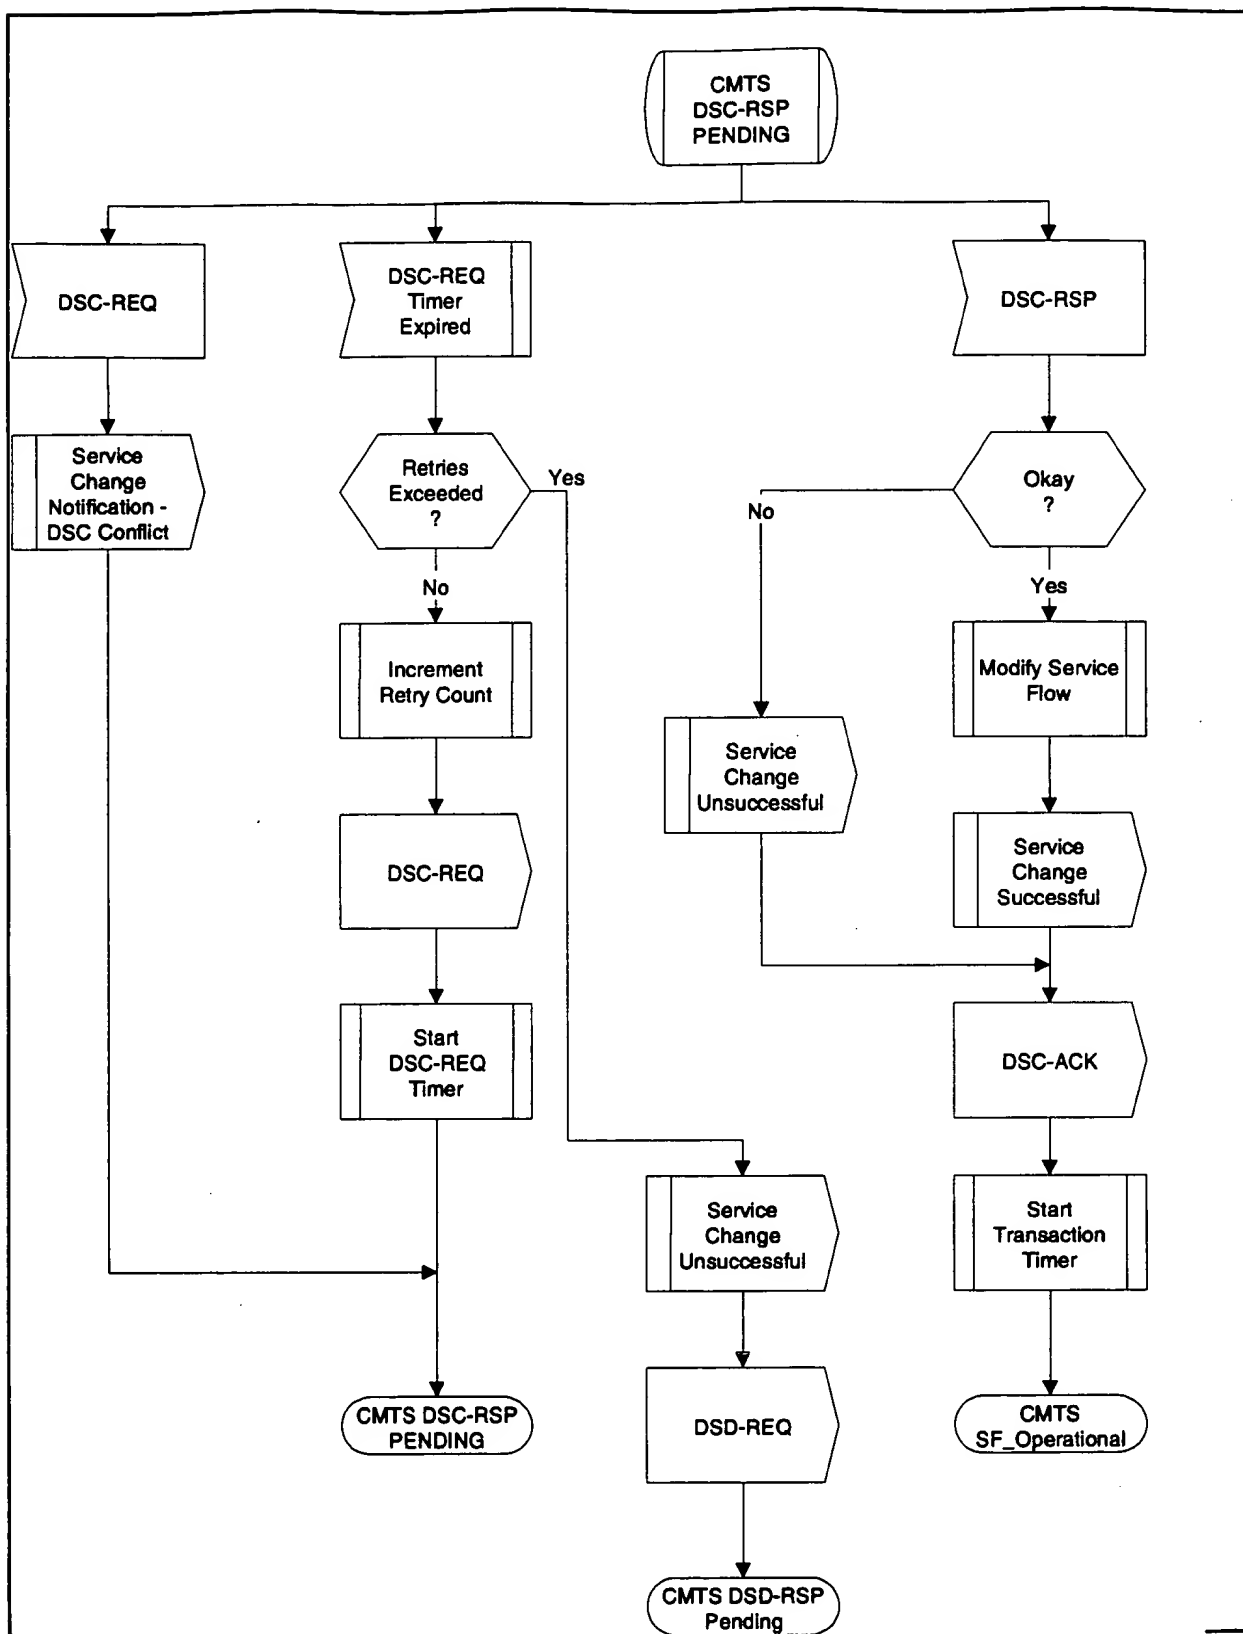


Figure 9-38. CMTS DSC-RSP Pending State (DSC Transactions) (figure edited 06/22/99 per rfi-n-99056. ew)

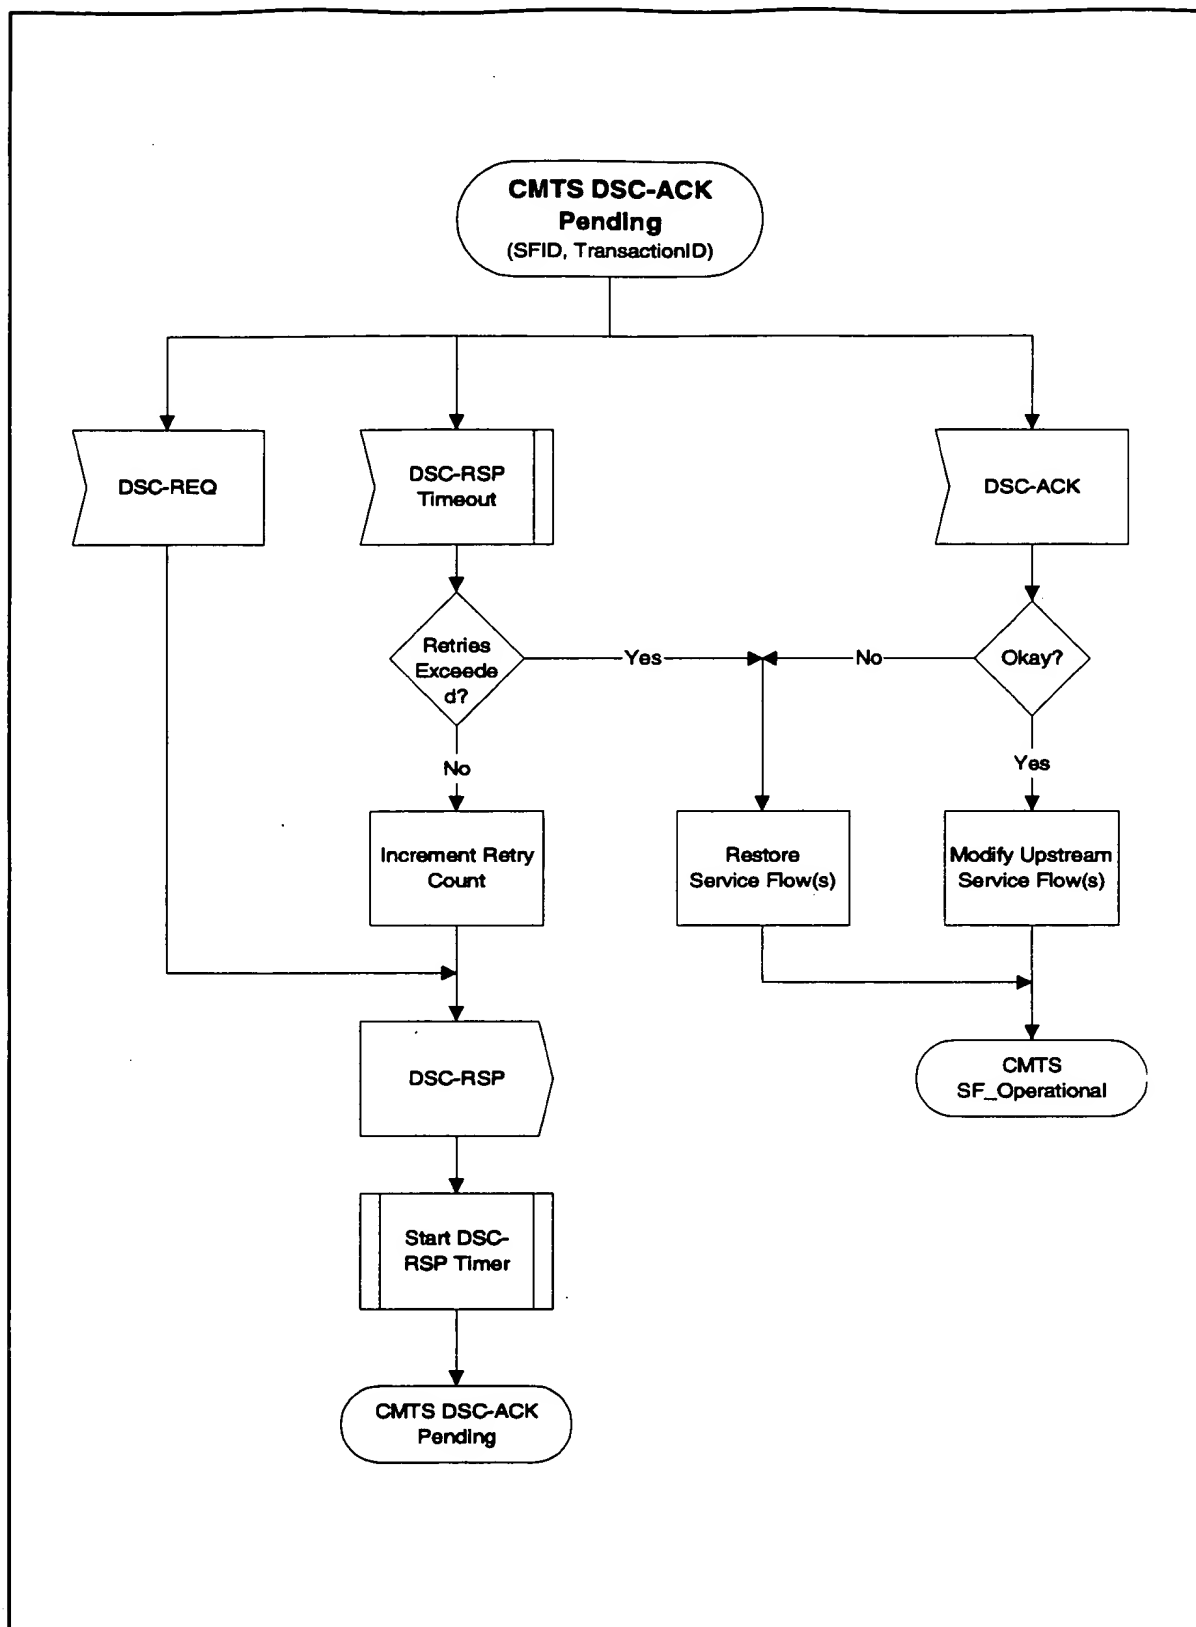


Figure 9-39. CMTS DSC-ACK Pending State (DSC Transactions)

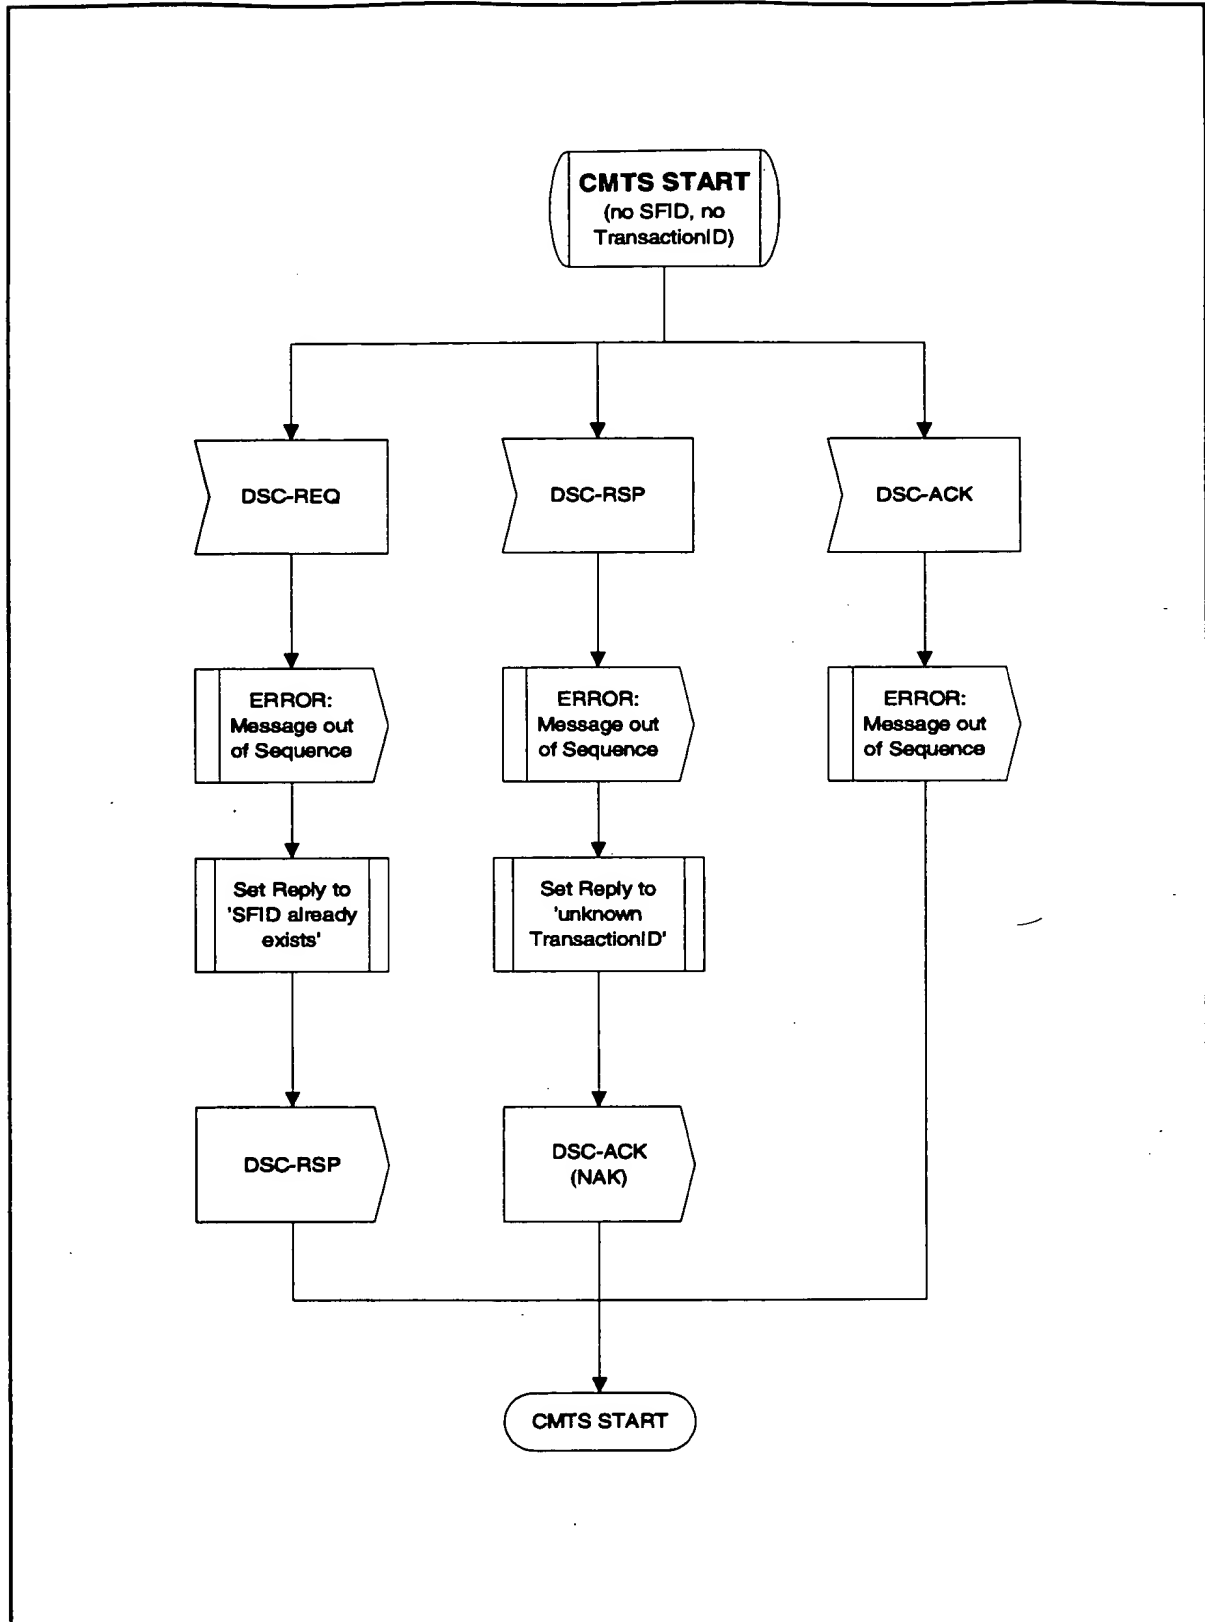


Figure 9-40. CMTS START State (DSC Transactions)

9.4.3 Dynamic Service Deletion

Any service flow can be deleted with the Dynamic Service Deletion (DSD) messages. When a Service Flow is deleted, all resources associated with it are released, including classifiers and PHS. However, if a Primary Service Flow of a CM is deleted, that CM is de-registered and **MUST** re-register. Also, if a Service Flow that was provisioned during registration is deleted, the provisioning information for that Service Flow is lost until the CM re-registers. However, the deletion of a provisioned Service Flow **MUST NOT** cause a CM to re-register. Therefore, care should be taken before deleting such Service Flows.

Note: Unlike DSA and DSC messages, DSD messages are limited to only a single Service Flow.

9.4.3.1 CM Initiated Dynamic Service Deletion

A CM wishing to delete a Service Flow generates a delete request to the CMTS using a Dynamic Service Deletion-Request message (DSD-REQ). The CMTS removes the Service Flow and generates a response using a Dynamic Service Deletion-Response message (DSD-RSP). Only one Service Flow can be deleted per DSD-Request.

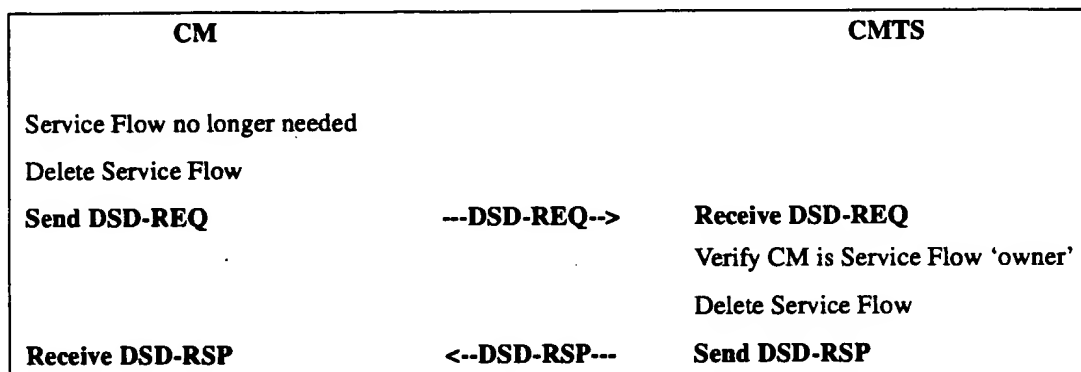


Figure 9-41. Dynamic Service Deletion Initiated from CM

9.4.3.2 CMTS Initiated Dynamic Service Deletion

A CMTS wishing to delete a dynamic Service Flow generates a delete request to the associated CM using a Dynamic Service Deletion-Request message (DSD-REQ). The CM removes the Service Flow and generates a response using a Dynamic Service Deletion-Response message (DSD-RSP). Only one Service Flow can be deleted per DSD-Request.

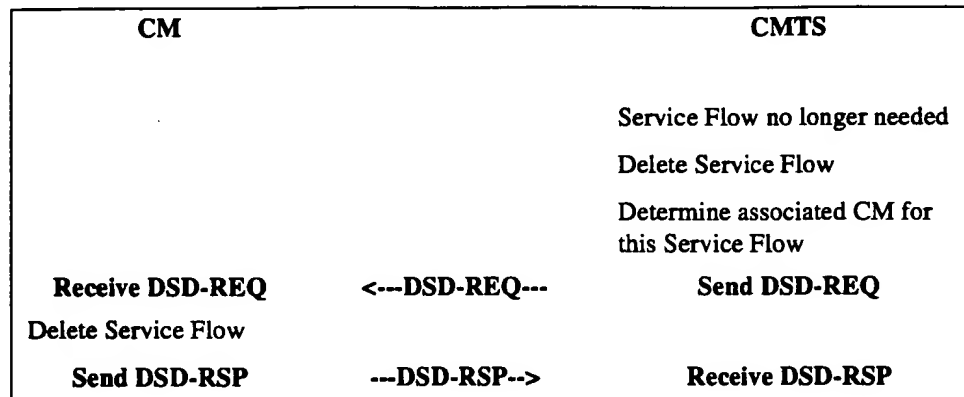


Figure 9-42. Dynamic Service Deletion Initiated from CMTS

9.4.3.3 Dynamic Service Deletion State Diagrams

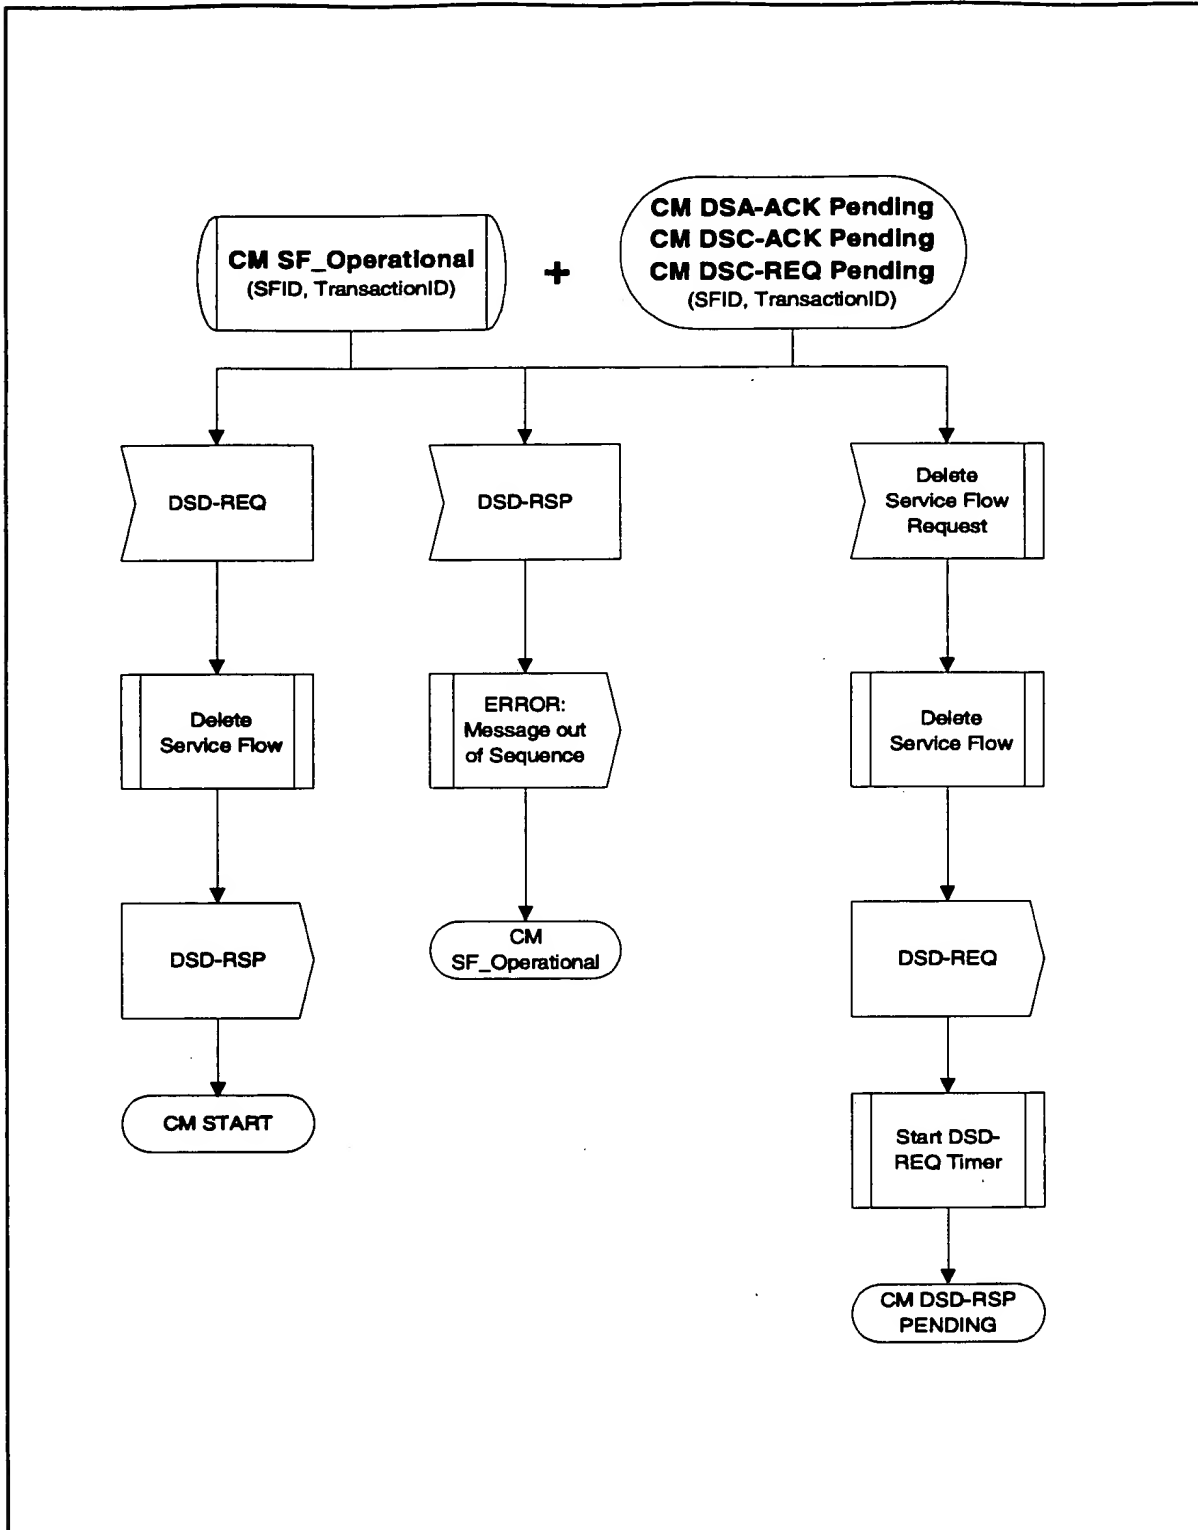


Figure 9-43. CM SF_Operational State (DSD Transactions)

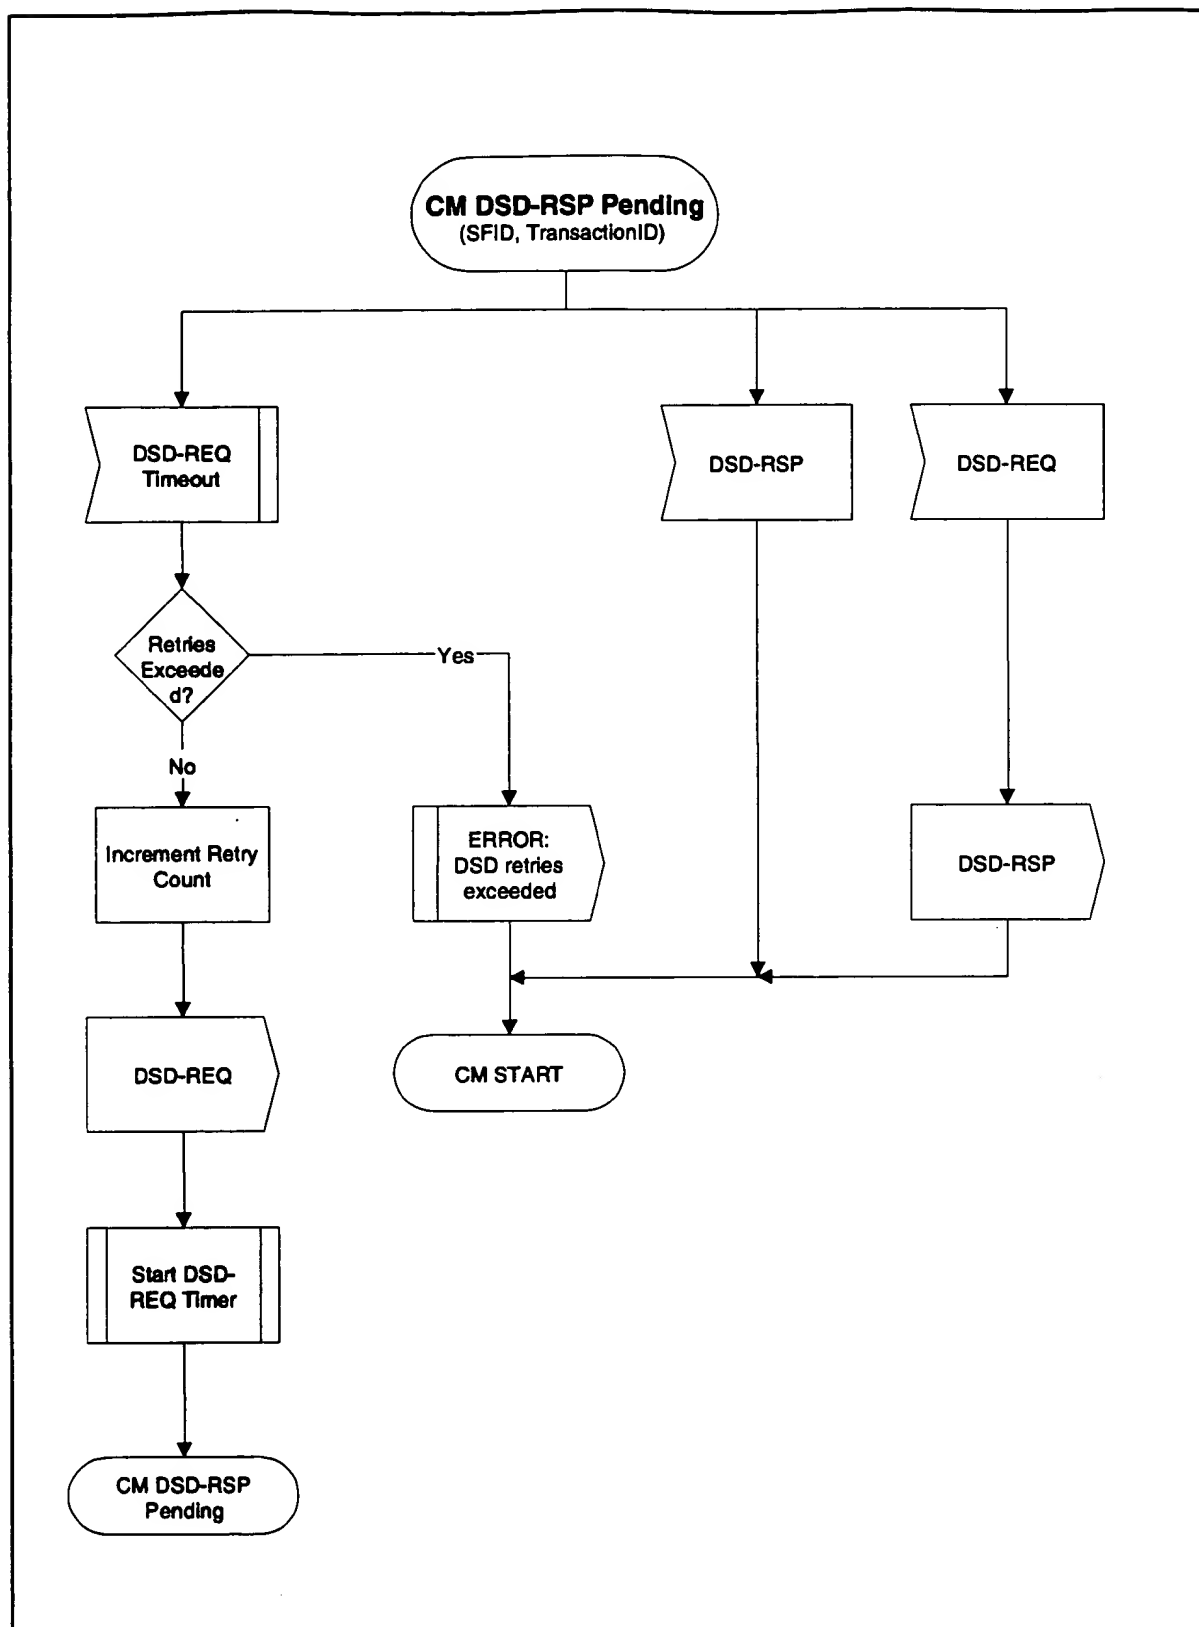


Figure 9-44. CM DSD-RSP Pending State (DSD Transactions) (figure edited 06/22/99 rfi-n-99043 ew)

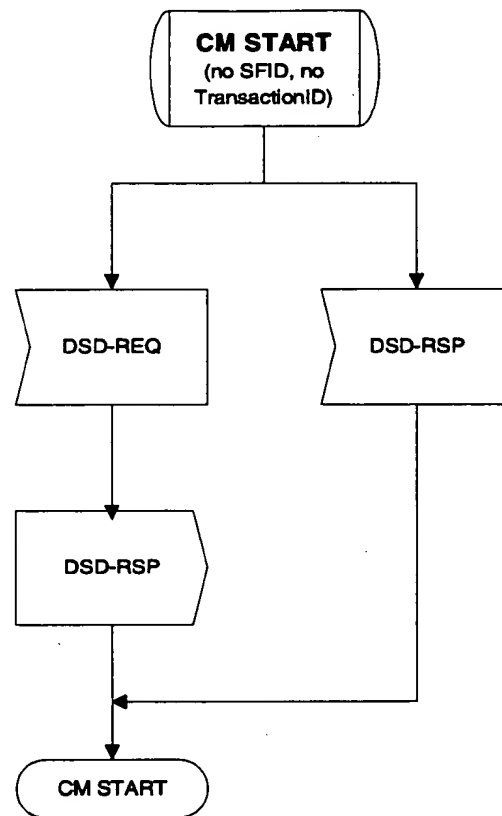


Figure 9-45. CM START State (DSD Transactions) (title edited 06/22/99 rfi-n-99043 ew)

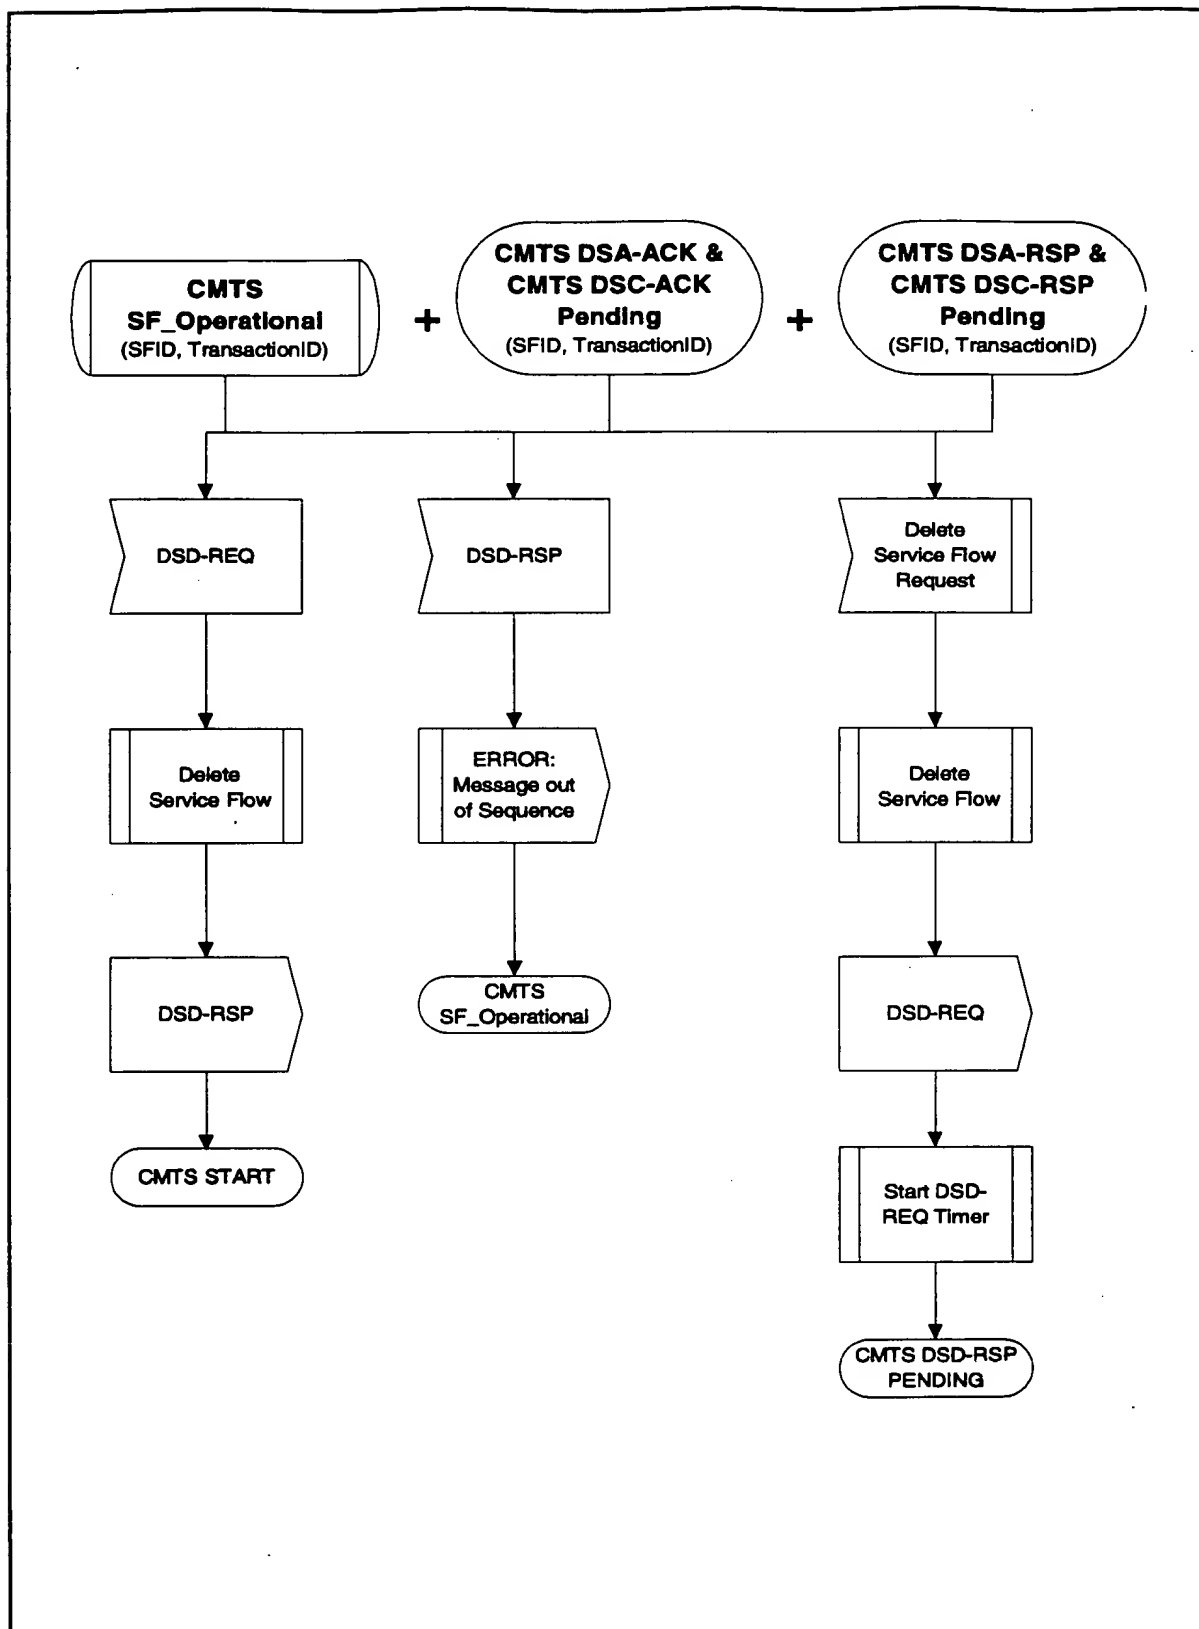


Figure 9-46. CMTS SF_Operational State (DSD Transactions)

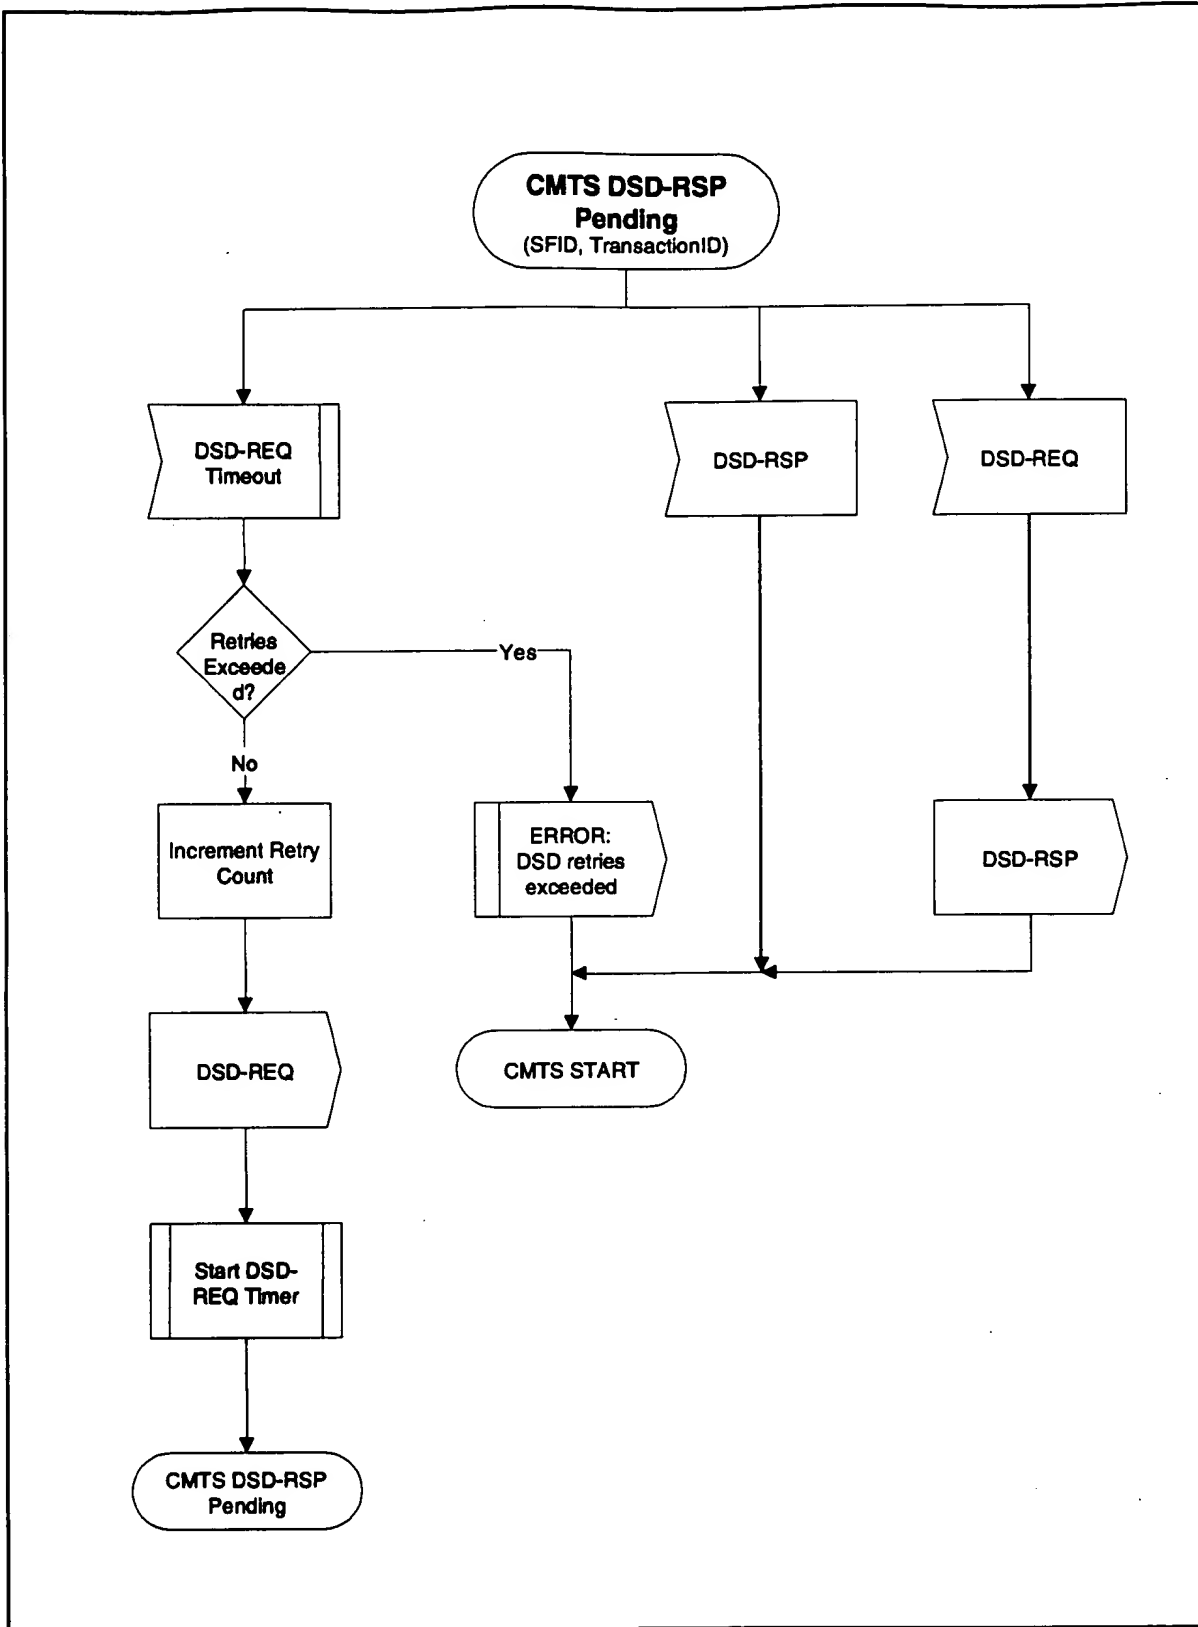


Figure 9-47. CMTS DSD-RSP Pending State (DSD Transactions) (figure edited 06/22/99 rfi-n-99043 ew)

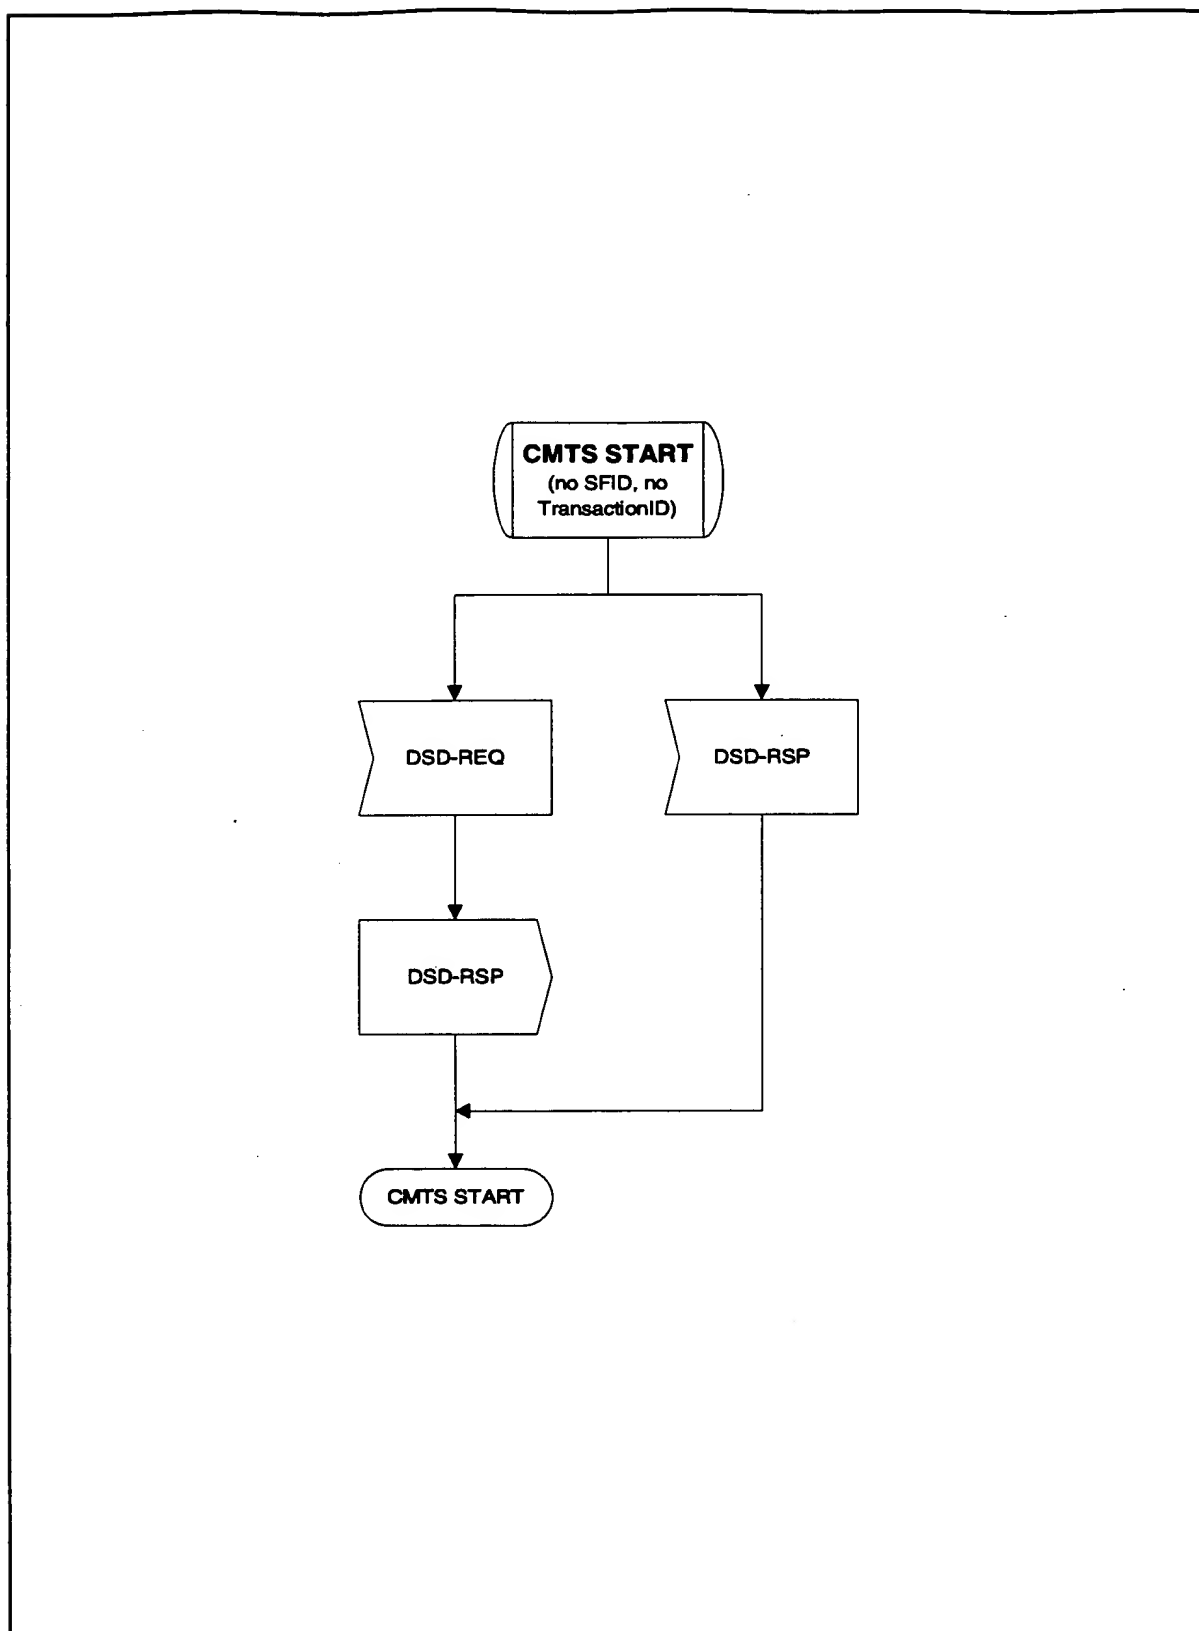


Figure 9-48. CMTS START State (DSD Transactions)

9.5 Fault Detection and Recovery

Fault detection and recovery occurs at multiple levels.

- At the physical level, FEC is used to correct errors where possible — refer to Section 4 for details.
- The MAC protocol protects against errors through the use of checksum fields across both the MAC Header and the data portions of the packet - refer to Section 6 for details.
- All MAC management messages are protected with a CRC covering the entire message, as defined in Section 6. Any message with a bad CRC **MUST** be discarded by the receiver.

Table 9-1 shows the recovery process that **MUST** be taken following the loss of a specific type of MAC message.

Appendix J contains a list of error codes with more useful information as to the failure of the PHY and MAC layers. Refer to Section 6.2.8 for additional information.

Table 9-1. Recovery Process on Loss of Specific MAC Messages

Message Name	Action Following Message Loss
SYNC	The CM can lose SYNC messages for a period of the Lost SYNC interval (see Appendix B) before it has lost synchronization with the network. A CM that has lost synchronization MUST NOT use the upstream and MUST try to re-establish synchronization.
UCD	A CM MUST receive a valid UCD before transmitting on the upstream. Failure to receive a valid UCD within the timeout period MUST cause the modem to reset and reinitialize its MAC connection.
MAP	A CM MUST NOT transmit without a valid upstream bandwidth allocation. If a MAP is missed due to error, the CM MUST NOT transmit for the period covered by the MAP.
RNG-REQ RNG-RSP	If a CM fails to receive a valid ranging response within a defined timeout period after transmitting a request, the request MUST be retried a number of times (as defined in Appendix B). Failure to receive a valid ranging response after the requisite number of attempts MUST cause the modem to reset and reinitialize its MAC connection.
REG-REQ REG-RSP	If a CM fails to receive a valid registration response within a defined timeout period after transmitting a request, the request will be retried a number of times (as defined in Appendix B). Failure to receive a valid registration response after the requisite number of attempts will cause the modem to reset and reinitialize its MAC connection.
UCC-REQ UCC-RSP	If a CMTS fails to receive a valid upstream channel change response within a defined timeout period after transmitting a request, the request MUST be retried a number of times (as defined in Appendix B). Failure to receive a valid response after the requisite number of attempts MUST cause the CMTS to consider the CM as unreachable.

Messages at the network layer and above are considered to be data packets by the MAC Sublayer. These are protected by the CRC field of the data packet and any packets with bad CRCs are discarded. Recovery from these lost packets is in accordance with the upper layer protocol.

9.5.1 Prevention of Unauthorized Transmissions

A CM **SHOULD** include a means for terminating RF transmission if it detects that its own carrier has been on continuously for longer than the longest possible valid transmission.

10 Supporting Future New Cable Modem Capabilities

10.1 Downloading Cable Modem Operating Software

A CMTS **SHOULD** be capable of being remotely reprogrammed in the field via a software download via the network.

The cable modem **MUST** be capable of being remotely reprogrammed in the field via a software download over the network. This software download capability **MUST** allow the functionality of the cable modem to be changed without requiring that cable system personnel physically revisit and reconfigure each unit. It is expected that this field programmability will be used to upgrade cable modem software to improve performance, accommodate new functions and features (such as enhanced class of service support), correct any design deficiencies discovered in the software, and to allow a migration path as the Data Over Cable Interface Specification evolves.

The mechanism used for download **MUST** be TFTP file transfer. The mechanism by which transfers are secured and authenticated is in [DOCSIS8]. The transfer **MUST** be initiated in one of two ways:

- An SNMP manager requests the CM to upgrade.
- If the Software Upgrade File Name in the CM's configuration file does not match the current software image of the CM, the CM **MUST** request the specified file via TFTP from the Software Server.

Note: The Software Server IP Address is a separate parameter. If present, the CM **MUST** attempt to download the specified file from this server. If not present, the CM **MUST** attempt to download the specified file from the configuration file server.

The CM **MUST** verify that the downloaded image is appropriate for itself. If the image is appropriate, the CM **MUST** write the new software image to non-volatile storage. Once the file transfer is completed successfully, the CM **MUST** restart itself with the new code image.

If the CM is unable to complete the file transfer for any reason, it **MUST** remain capable of accepting new software downloads (without operator or user interaction), even if power or connectivity is interrupted between attempts. The CM **MUST** log the failure and **MAY** report it asynchronously to the network manager.

Following upgrade of the operational software, the CM **MAY** need to follow one of the procedures described above in order to change channels to use the enhanced functionality.

If the CM is to continue to operate in the same upstream and downstream channels as before the upgrade, then it **MUST** be capable of inter-working with other CMs which **MAY** be running previous releases of software.

Where software has been upgraded to meet a new version of the specification, then it is critical that it **MUST** inter-work with the previous version in order to allow a gradual transition of units on the network.

This page intentionally left blank.

Appendix A. Well-Known Addresses

A.1 MAC Addresses

MAC addresses described here are defined using the Ethernet/ISO8802-3 convention as bit-little-endian.

The following multicast address **MUST** be used to address the set of all CM MAC sublayers; for example, when transmitting Allocation Map PDUs.

01-E0-2F-00-00-01

The address range

01-E0-2F-00-00-03 through 01-E0-2F-00-00-0F

is reserved for future definition. Frames addressed to any of these addresses **SHOULD NOT** be forwarded out of the MAC-sublayer domain.

A.2 MAC Service IDs

The following MAC Service IDs have assigned meanings. Those not included in this table are available for assignment, either by the CMTS or administratively.

A.2.1 All CMs and No CM Service IDs

These Service IDs are used in MAPs for special purposes or to indicate that any CM can respond in the corresponding interval.

0x0000	Addressed to no CM. Typically used when changing upstream burst parameters so that CMs have time to adjust their modulators before the new upstream settings are in effect.
0x3FFF	Addressed to all CMs. Typically used for broadcast Request intervals or Initial Maintenance intervals.

A.2.2 Well-Known 'Multicast' Service IDs

These Service IDs are only used for Request/Data IE's. They indicate that any CM can respond in a given interval, but that it must limit the size of its transmission to a particular number of minislots (as indicated by the particular multicast SID assigned to the interval).

0x3FF1-0x3FFE	Addressed to all CMs. Available for small data PDUs, as well as requests (used only with request/data IEs). The last digit indicates the frame length and transmission opportunities as follows:
0x3FF1	Within the interval specified, a transmission may start at any mini-slot, and must fit within one mini-slot.
0x3FF2	Within the interval specified, a transmission may start at every other mini-slot, and must fit within two mini-slots (e.g., a station may start transmission on the first mini-slot within the interval, the third mini-slot, the fifth, etc.).
0x3FF3	Within the interval specified, a transmission may start at any third mini-slot, and must fit within three mini-slots (e.g., starts at first, fourth, seventh, etc.).

0x3FF4	Starts at first, fifth, ninth, etc.
...	
0x3FFD	Starts at first, fourteenth (14 th), twenty-seventh (27 th), etc.
0x3FFE	Within the interval specified, a transmission may start at any 14 th mini-slot, and must fit within 14 mini-slots.

A.2.3 Priority Request Service IDs

These Service IDs (0x3Exx) are reserved for Request IEs (refer to C.2.2.5.2).

- If 0x01 bit is set, priority zero can request
- If 0x02 bit is set, priority one can request
- If 0x04 bit is set, priority two can request
- If 0x08 bit is set, priority three can request
- If 0x10 bit is set, priority four can request
- If 0x20 bit is set, priority five can request
- If 0x40 bit is set, priority six can request
- If 0x80 bit is set, priority seven can request

Bits can be combined as desired by the CMTS upstream scheduler for any Request IUCs.

A.3 MPEG PID

All DOCSIS data **MUST** be carried in MPEG-2 packets with the header PID field set to 0x1FFE.

Appendix B. Parameters and Constants

System	Name	Time Reference	Minimum Value	Default Value	Maximum Value
CMTS	Sync Interval	Nominal time between transmission of SYNC messages (ref 6.3.2)			200 msec
CMTS	UCD Interval	Time between transmission of UCD messages (ref. 6.3.3)			2 sec
CMTS	Max MAP Pending	The number of mini-slots that a CMTS is allowed to map into the future (ref. 6.3.4)			4096 mini-slot times
CMTS	Ranging Interval	Time between transmission of broadcast Ranging requests (ref. 7.3.3)			2 sec
CM	Lost Sync Interval	Time since last received Sync message before synchronization is considered lost			600 msec
CM	Contention Ranging Retries	Number of Retries on contention Ranging Requests (ref. 9.2.4)	16		
CM, CMTS	Invited Ranging Retries	Number of Retries on Inviting Ranging Requests (ref. 9.2.4)	16		
CM	Request Retries	Number of retries on bandwidth allocation requests	16		
CM	Registration Request Retries	Number of retries on registration requests	3		
CM	Data Retries	Number of retries on immediate data transmission	16		
CMTS	CM MAP processing time	Time provided between arrival of the last bit of a MAP at a CM and effectiveness of that MAP (ref. 7.1.1)	200 μ s		
CMTS	CM Ranging Response processing time	Minimum time allowed for a CM following receipt of a ranging response before it is expected to reply to an invited ranging request	1 msec		
CMTS	CM Configuration	The maximum time allowed for a CM, following receipt of a configuration file, to send a Registration Request to a CMTS.	30 sec		
CM	T1	Wait for UCD timeout			5 * UCD interval maximum value
CM	T2	Wait for broadcast ranging timeout			5 * ranging interval
CM	T3	Wait for ranging response	50 msec	200 msec	200 msec
CM	T4	Wait for unicast ranging opportunity. If the pending-till-complete field was used earlier by this modem, then the value of that field must be added to this interval.	30 sec		35 sec
CMTS	T5	Wait for Upstream Channel Change response			2 sec
CM	T6	Wait for registration response			3 sec
CM CMTS	Mini-slot size	Size of mini-slot for upstream transmission. Must be a power of 2 (in units of the Timebase Tick)	32 symbol times		
CM CMTS	Timebase Tick	System timing unit	6.25 μ sec		

CM CMTS	DSx Request Retries	Number of Retries on DSA/DSC/DSD Requests	3		
CM CMTS	DSx Response Retries	Number of Retries on DSA/DSC/DSD Responses	3		
CM CMTS	T7	Wait for DSA/DSC/DSD Response timeout			3 sec
CM CMTS	T8	Wait for DSA/DSC Acknowledge timeout			300 msec
CM	TFTP Backoff Start	Initial value for TFTP backoff	1sec		
CM	TFTP Backoff End	Last value for TFTP backoff	16 sec		
CM	TFTP Request Retries	Number of retries on TFTP request	16		
CM	TFTP Download Retries	Number of retries on entire TFTP downloads	3		
CM	TFTP Wait	The wait between TFTP retry sequences	10 min		
CM	ToD Retries	Number of Retries per ToD Retry Period	3		
CM	ToD Retry Period	Time period for ToD retries	5 min		
CMTS	T9	Registration Timeout, the time allowed between the CMTS sending a RNG-RSP (success) to a CM, and receiving a REG-REQ from that same CM.	15 min	15 min ^a	

a. Row added per rfi-n-99054 06/29/99. ew

Appendix C. Common Radio Frequency Interface Encodings

C.1 Encodings for Configuration and MAC-Layer Messaging

The following type/length/value encodings **MUST** be used in both the configuration file (see Appendix D.), in CM registration requests and in Dynamic Service Messages. All multi-octet quantities are in network-byte order, i.e., the octet containing the most-significant bits is the first transmitted on the wire.

The following configuration settings **MUST** be supported by all CMs which are compliant with this specification.

C.1.1 Configuration File and Registration Settings

These settings are found in the configuration file and, if present, **MUST** be forwarded by the CM to the CMTS in its Registration Request.

C.1.1.1 Downstream Frequency Configuration Setting

The receive frequency to be used by the CM. It is an override for the channel selected during scanning. This is the center frequency of the downstream channel in Hz stored as a 32-bit binary number.

Type	Length	Value
1	4	Rx Frequency

Valid Range:

The receive frequency **MUST** be a multiple of 62500 Hz.

C.1.1.2 Upstream Channel ID Configuration Setting

The upstream channel ID which the CM **MUST** use. The CM **MUST** listen on the defined downstream channel until an upstream channel description message with this ID is found. It is an override for the channel selected during initialization.

Type	Length	Value
2	1	Channel ID

C.1.1.3 Network Access Control Object

If the value field is a 1, CPE attached to this CM are allowed access to the network, based on CM provisioning. If the value of this field is a 0, the CM **MUST NOT** forward traffic from attached CPE to the RF MAC network, but **MUST** continue to accept and generate traffic from the CM itself. The value of this field does not affect CMTS service flow operation and does not affect CMTS data forwarding operation.¹

Type	Length	On / Off
3	1	1 or 0

1. Paragraph edited per rfi-n-99052 06/29/99. ew

Note: The intent of "NACO = 0" is that the CM does not forward traffic from any attached CPE onto the cable network. (A CPE is any client device attached to that CM, regardless of how that attachment is implemented.) However, with "NACO = 0", management traffic to the CM is not restricted. Specifically, with NACO off, the CM remains manageable, including sending/receiving management traffic such as (but not limited to):

- ARP: allow the modem to resolve IP addresses, so it can respond to queries or send traps.
- DHCP: allow the modem to renew its IP address lease.
- ICMP: enable network troubleshooting for tools such as "ping" and "traceroute."
- ToD: allow the modem to continue to synchronize its clock after boot.
- TFTP: allow the modem to download either a new configuration file or a new software image.
- SYSLOG: allow the modem to report network events.
- SNMP: allow management activity

In DOCSIS v1.1, with NACO off, the primary upstream and primary downstream service flows of the CM remain operational only for management traffic to and from the CM. With respect to DOCSIS v1.1 provisioning, a CMTS should ignore the NACO value and allocate any service flows that have been authorized by the provisioning server.¹

1. section added per rfi-n-99052 06/28/99. ew

C.1.1.4 DOCSIS 1.0 Class of Service Configuration Setting

This field defines the parameters associated with a DOCSIS 1.0 class of service. Any CM registering with a DOCSIS 1.0 Class of Service Configuration Setting will be treated as a DOCSIS 1.0 CM. Refer to Section 6.3.8.

This field defines the parameters associated with a class of service. It is somewhat complex in that it is composed from a number of encapsulated type/length/value fields. The encapsulated fields define the particular class of service parameters for the class of service in question. Note that the type fields defined are only valid within the encapsulated class of service configuration setting string. A single class of service configuration setting is used to define the parameters for a single service class. Multiple class definitions use multiple class of service configuration setting sets.

Type	Length	Value
4	n	

C.1.1.4.1 Class ID

The value of the field specifies the identifier for the class of service to which the encapsulated string applies.

Type	Length	Value
4.1	1	

Valid Range

The class ID MUST be in the range 1 to 16.

C.1.1.4.2 Maximum Downstream Rate Configuration Setting

For a single SID modem, the value of this field specifies the maximum downstream rate in bits per second that the CMTS is permitted to forward to CPE unicast MAC addresses learned or configured as mapping to the registering modem.

For a multiple SID modem, the aggregate value of these fields specifies the maximum downstream rate in bits per second that the CMTS is permitted to forward to CPE unicast MAC addresses learned or configured as mapping to the registering modem.

This is the peak data rate for Packet PDU Data (including destination MAC address and the CRC) over a one-second interval. This does not include MAC packets addressed to broadcast or multicast MAC addresses. The CMTS MUST limit downstream forwarding to this rate. The CMTS MAY delay, rather than drop, over-limit packets.

Type	Length	Value
4.2	4	

Note: This is a limit, not a guarantee that this rate is available.

C.1.1.4.3 Maximum Upstream Rate Configuration Setting

The value of this field specifies the maximum upstream rate in bits per second that the CM is permitted to forward to the RF Network.

This is the peak data rate for Packet PDU Data (including destination address and the CRC) over a one-second interval. The CM MUST limit all upstream forwarding (both contention and reservation-based), for the corresponding SID, to this rate. The CM MUST include Packet PDU Data packets addressed to broadcast or multicast addresses when calculating this rate.

The CM MUST enforce the maximum upstream rate. It SHOULD NOT discard upstream traffic simply because it exceeds this rate.

The CMTS MUST enforce this limit on all upstream data transmissions, including data sent in contention. The CMTS SHOULD generate an alarm if a modem exceeds its allowable rate.

Type	Length	Value
4.3	4	

Note:The purpose of this parameter is for the CM to perform traffic shaping at the input to the RF network and for the CMTS to perform traffic policing to ensure that the CM does not exceed this limit.

The CMTS could enforce this limit by any of the following methods:

- a) discarding over-limit requests.
- b) deferring (through zero-length grants) the grant until it is conforming to the allowed limit.
- c) discarding over-limit data packets.
- d) Reporting to a policy monitor (for example, using the alarm mechanism) that is capable of incapacitating errant CMs.

Note:This is a limit, not a guarantee that this rate is available.

C.1.1.4.4 Upstream Channel Priority Configuration Setting

The value of the field specifies the relative priority assigned to this service class for data transmission in the upstream channel. Higher numbers indicate higher priority.

Type	Length	Value
4.4	1	

Valid Range
0 -> 7

C.1.1.4.5 Guaranteed Minimum Upstream Channel Data Rate Configuration Setting

The value of the field specifies the data rate in bit/sec which will be guaranteed to this service class on the upstream channel.

Type	Length	Value
4.5	4	

C.1.1.4.6 Maximum Upstream Channel Transmit Burst Configuration Setting

The value of the field specifies the maximum transmit burst (in bytes) which this service class is allowed on the upstream channel. A value of zero means there is no limit. Note: This value does not include any physical layer overhead.

Type	Length	Value
4.6	2	

C.1.1.4.7 Class-of-Service Privacy Enable

This configuration setting enables/disables Baseline Privacy on a provisioned CoS. See [DOCSIS8].

Type	Length	Enable / Disable
4.7 (= CoS_BP_ENABLE)	1	1 or 0

Table C-1. Sample DOCSIS 1.0 Class of Service Encoding

Type	Length	Value (sub)type	Length	Value	
4	28				class of service configuration setting
		1	1	1	service class 1
		2	4	10,000,000	max. downstream rate of 10 Mb/sec
		3	4	300,000	max. upstream rate of 300 kbps
		4	1	5	return path priority of 5
		5	4	64000	min guaranteed 64 kb/sec
4	28	6	2	1518	max. Tx burst of 1518 bytes
					class of service configuration setting
		1	1	2	service class 2
		2	4	5,000,000	max. forward rate of 5 Mb/sec
		3	4	300,000	max. return rate of 1 Mb/sec
		4	1	3	return path priority of 3
		5	4	32000	min guaranteed 32 kb/sec
		6	2	1518	max. Tx burst of 1518 bytes

C.1.1.5 CM Message Integrity Check (MIC) Configuration Setting

The value field contains the CM message integrity check code. This is used to detect unauthorized modification or corruption of the configuration file.

Type	Length	Value
6	16	d1 d2..... d16

C.1.1.6 CMTS Message Integrity Check (MIC) Configuration Setting

The value field contains the CMTS message integrity check code. This is used to detect unauthorized modification or corruption of the configuration file.

Type	Length	Value
7	16	d1 d2..... d16

C.1.1.7 Maximum Number of CPEs

The maximum number of CPEs that can be granted access through a CM during a CM epoch. The CM epoch is (from Section 3.1.2.3.1) the time between startup and hard reset of the modem. The maximum number of CPE's MUST be enforced by the CM.

Note: This parameter should not be confused with the number of CPE addresses a CM may learn. A modem may learn Ethernet MAC addresses up to its maximum number of CPE addresses (from Section 3.1.2.3.1). The maximum number of CPEs that are granted access through the modem is governed by this configuration setting.

Type	Length	Value
18	1	

If present, the value MUST be positive and non-zero. The non-existence of this option means the default value of 1.

Note: This is a limit on the maximum number of CPEs a CM will grant access to. Hardware limitations of a given modem implementation may require the modem to use a lower value.

C.1.1.8 TFTP Server Timestamp

The sending time of the configuration file in seconds. The definition of time is as in [RFC-868]

Type	Length	Value
19	4	Number of seconds since 00:00 1 Jan 1900

Note: The purpose of this parameter is to prevent replay attacks with old configuration files.

C.1.1.9 TFTP Server Provisioned Modem Address

The IP Address of the modem requesting the configuration file.

Type	Length	Value
20	4	IP Address

Note: The purpose of this parameter is to prevent IP spoofing during registration.

C.1.1.10 Upstream Packet Classification Configuration Setting

This field defines the parameters associated with one entry in an upstream traffic classification list. Refer to Section C.2.1.1.

Type	Length	Value
22	n	

C.1.1.11 Downstream Packet Classification Configuration Setting

This field defines the parameters associated with one Classifier in an downstream traffic classification list. Refer to Section C.2.1.2.

Type	Length	Value
23	n	

C.1.1.12 Upstream Service Flow Encodings

This field defines the parameters associated with upstream scheduling for one Service Flow. Refer to Section C.2.2.1.

Type	Length	Value
24	n	

C.1.1.13 Downstream Service Flow Encodings

This field defines the parameters associated with downstream scheduling for one Service Flow. Refer to Section C.2.2.2.

Type	Length	Value
25	n	

C.1.1.14 Payload Header Suppression

This field defines the parameters associated with Payload Header Suppression.

Type	Length	Value
26	n	

C.1.1.15 Maximum Number of Classifiers

This is the maximum number of Classifiers that the CM is allowed to have active.

This is necessary when using deferred activation since the number of provisioned Service Flows may be high and since each Service Flow might support multiple Classifiers. Provisioning represents the set of Service Flows the CM can choose between, however, it may still be desirable to limit the number of simultaneously admitted Classifiers applied to this set. This parameter provides the ability to limit the size of that set.

Type	Length	Value
28	2	Maximum number of simultaneous admitted classifiers

The default value is 0 — no limit.

C.1.1.16 Privacy Enable

This configuration setting enables/disables Baseline Privacy on the Primary Service Flow and all other Service Flows for this CM.

Type	Length	Value
29	1	0 — Disable 1 — Enable

The default value of this parameter is 1 — privacy enabled.

C.1.1.17 Vendor-Specific Information

Vendor-specific information for cable modems, if present, **MUST** be encoded in the vendor specific information field (VSIF) (code 43) using the Vendor ID field (C.1.3.2) to specify which TLV tuples apply to which vendors products. The Vendor ID **MUST** be the first TLV embedded inside VSIF. If the first TLV inside VSIF is not a Vendor ID, then the TLV must be discarded.

This configuration setting **MAY** appear multiple times. The same Vendor ID **MAY** appear multiple times. This configuration setting **MAY** be nested inside a Packet Classification Configuration Setting, a Service Flow Configuration Setting, or a Service Flow Response. However, there **MUST NOT** be more than one Vendor ID TLV inside a single VSIF.

Type	Length	Value
43	n	per vendor definition

Example:

Configuration with vendor A specific fields and vendor B specific fields:

VSIF (43) + n (number of bytes inside this VSIF)
8 (Vendor ID Type) + 3 (length field) + Vendor ID of Vendor A
Vendor A Specific Type #1 + length of the field + Value #1
Vendor A Specific Type #2 + length of the field + Value #2

VSIF (43) + m (number of bytes inside this VSIF)
8 (Vendor ID Type) + 3 (length field) + Vendor ID of Vendor B
Vendor B Specific Type + length of the field + Value

C.1.2 Configuration-File-Specific Settings

These settings are found in only the configuration file. They **MUST NOT** be forwarded to the CMTS in the Registration Request.

C.1.2.1 End-of-Data Marker

This is a special marker for end of data.

It has no length or value fields.

Type
255

C.1.2.2 Pad Configuration Setting

This has no length or value fields and is only used following the end of data marker to pad the file to an integral number of 32-bit words.

Type
0

C.1.2.3 Software Upgrade Filename

The filename of the software upgrade file for the CM. The filename is a fully qualified directory-path name. The file is expected to reside on a TFTP server identified in a configuration setting option defined in Appendix D.1.1. See Section 10.1.

Type	Length	Value
9	n	filename

C.1.2.4 SNMP Write-Access Control

This object makes it possible to disable SNMP "Set" access to individual MIB objects. Each instance of this object controls access to all of the writeable MIB objects whose Object ID (OID) prefix matches. This object may be repeated to disable access to any number of MIB objects.

Type	Length	Value
10	n	OID prefix plus control flag

Where n is the size of the ASN.1 Basic Encoding Rules [ISO8025] encoding of the OID prefix plus one byte for the control flag.

The control flag may take values:

- 0 - allow write-access
- 1 - disallow write-access

Any OID prefix may be used. The Null OID 0.0 may be used to control access to all MIB objects. (The OID 1.3.6.1 will have the same effect.)

When multiple instances of this object are present and overlap, the longest (most specific) prefix has precedence. Thus, one example might be

someTable	disallow write-access
someTable.1.3	allow write-access

This example disallows access to all objects in someTable except for someTable.1.3.

C.1.2.5 SNMP MIB Object

This object allows arbitrary SNMP MIB objects to be Set via the TFTP-Registration process.

Type	Length	Value
11	n	variable binding

where the value is an SNMP VarBind as defined in [RFC-1157]. The VarBind is encoded in ASN.1 Basic Encoding Rules, just as it would be if part of an SNMP Set request.

The cable modem **MUST** treat this object as if it were part of an SNMP Set Request with the following caveats:

- It **MUST** treat the request as fully authorized (it cannot refuse the request for lack of privilege).
- SNMP Write-Control provisions (see previous section) do not apply.
- No SNMP response is generated by the CM.

This object **MAY** be repeated with different VarBinds to “Set” a number of MIB objects. All such Sets **MUST** be treated as if simultaneous.

Each VarBind **MUST** be limited to 255 bytes.

C.1.2.6 CPE Ethernet MAC Address

This object configures the CM with the Ethernet MAC address of a CPE device (see Section 3.1.2.3.1). This object may be repeated to configure any number of CPE device addresses.

Type	Length	Value
14	6	Ethernet MAC Address of CPE

C.1.2.7 Software Upgrade TFTP Server

The IP address of the TFTP server, on which the software upgrade file for the CM resides. See Section 10.1 and Appendix C.1.2.3

Type	Length	Value
21	4	ip1,ip2,ip3,ip4

C.1.3 Registration-Request/Response-Specific Encodings

These encodings are not found in the configuration file, but are included in the Registration Request. Some encodings are also used in the Registration Response.

The CM **MUST** include Modem Capabilities Encodings in its Registration Request. If present in the corresponding Registration Request, the CMTS **MUST** include Modem Capabilities in the Registration Response.

C.1.3.1 Modem Capabilities Encoding

The value field describes the capabilities of a particular modem, i.e., implementation dependent limits on the particular features or number of features which the modem can support. It is composed from a number of encapsulated type/length/value fields. The encapsulated sub-types define the specific capabilities for the modem in question. Note that the sub-type fields defined are only valid within the encapsulated capabilities configuration setting string.

Type	Length	Value
5	n	

The set of possible encapsulated fields is described below.

C.1.3.1.1 Concatenation Support

If the value field is a 1 the CM requests concatenation support from the CMTS.

Type	Length	On / Off
5.1	1	1 or 0

C.1.3.1.2 DOCSIS Version

DOCSIS version of this modem.

Type	Length	Value
5.2	1	0: DOCSIS v1.0 1: DOCSIS v1.1 2-255: Reserved

If this tuple is absent, the CMTS **MUST** assume DOCSIS v1.0 operation. The absence of this tuple or the value 'DOCSIS 1.0' does not necessarily mean the CM only supports DOCSIS 1.0 functionality — the CM **MAY** indicate it supports other individual capabilities with other Modem Capability Encodings. (Refer to G.3)

C.1.3.1.3 Fragmentation Support

If the value field is a 1 the CM requests fragmentation support from the CMTS.

Type	Length	Value
5.3	1	1 or 0

C.1.3.1.4 Payload Header Suppression Support

If the value field is a 1 the CM requests payload header suppression support from the CMTS.

Type	Length	Value
5.4	1	1 or 0

C.1.3.1.5 IGMP Support

If the value field is a 1 the CM supports DOCSIS 1.1-compliant IGMP.

Type	Length	Value
5.5	1	1 or 0

C.1.3.1.6 Privacy Support

The value is the BPI support of the CM.

Type	Length	Value
5.6	1	0 BPI Support
		1 BPI Plus Support
		2 - 255 Reserved

C.1.3.1.7 Downstream SAID Support

The field shows the number of Downstream SAIDs the modem can support.

Type	Length	Value
5.7	1	Number of Downstream SAIDs the CM can support.

If the number of SAIDs is 0 that means the Modem can support only 1 SAID.

C.1.3.1.8 Upstream SID Support

The field shows the number of Upstream SIDs the modem can support.

Type	Length	Value
5.8	1	Number of Upstream SIDs the CM can support.

If the number of SIDs is 0 that means the Modem can support only 1 SID.

C.1.3.1.9 Optional Filtering Support

The fields shows the optional filtering support in the modem.

Type	Length	Value
5.9	1	Packet Filtering Support Array
		bit #0: 802.1P filtering
		bit #1: 802.1Q filtering
		bit #2-7: reserved must be set to zero

C.1.3.2 Vendor ID Encoding

The value field contains the vendor identification specified by the three-byte vendor-specific Organization Unique Identifier of the CM MAC address.

The Vendor ID **MUST** be used in a Registration Request, but **MUST NOT** be used as a stand-alone configuration file element. It **MAY** be used as a sub-field of the Vendor Specific Information Field in a configuration file. When used as a sub-field of the Vendor Specific Information field, this identifies the Vendor ID of the CMs which are intended to use this information. When the vendor ID is used in a Registration Request, then it is the Vendor ID of the CM sending the request.

Type	Length	Value
8	3	v1, v2, v3

C.1.3.3 Modem IP Address

For backwards compatibility with DOCSIS v 1.0. Replaced by 'TFTP Server Provisioned Modem Address'.

Type	Length	Value
12	4	IP Address

C.1.3.4 Service(s) Not Available Response

This configuration setting **MUST** be included in the Registration Response message if the CMTS is unable or unwilling to grant any of the requested classes of service that appeared in the Registration Request. Although the value applies only to the failed service class, the entire Registration Request **MUST** be considered to have failed (none of the class-of-service configuration settings are granted).

Type	Length	Value
13	3	Class ID, Type, Confirmation Code

Where

Class ID	is the class-of-service class from the request which is not available
Type	is the specific class-of-service object within the class which caused the request to be rejected
Confirmation Code	Refer to C.4.

C.1.4 Dynamic-Service-Message-Specific Encodings

These encodings are not found in the configuration file, nor in the Registration Request/Response signaling. They are only found in Dynamic Service Addition, Dynamic Service Change and Dynamic Service Deletion Request/Response messages:

C.1.4.1 HMAC-Digest

The HMAC-Digest setting is a keyed message digest. If privacy is enabled, the HMAC-Digest Attribute **MUST** be the final Attribute in the Dynamic Service message's Attribute list. The message digest is performed over the all of the Dynamic Service parameters (starting immediately after the MAC Management Message Header and up to, but not including the HMAC Digest setting), other than the HMAC-Digest, in the order in which they appear within the packet.

Inclusion of the keyed digest allows the receiver to authenticate the message. The HMAC-Digest algorithm, and the upstream and downstream key generation requirements are documented in [DOCSIS-BPI+].

This parameter contains a keyed hash used for message authentication. The HMAC algorithm is defined in [RFC2104]. The HMAC algorithm is specified using a generic cryptographic hash algorithm. Baseline Privacy uses a particular version of HMAC that employs the Secure Hash Algorithm (SHA-1), defined in [SHA].

A summary of the HMAC-Digest Attribute format is shown below. The fields are transmitted from left to right.

Type	Length	Value
27	20	A 160-bit (20 octet) keyed SHA hash

C.2 Quality-of-Service-Related Encodings

C.2.1 Packet Classification Encodings

The following type/length/value encodings **MUST** be used in both the configuration file, registration messages, and Dynamic Service messages to encode parameters for packet classification and scheduling. All multi-octet quantities are in network-byte order, i.e., the octet containing the most-significant bits is the first transmitted on the wire.

The following configuration settings **MUST** be supported by all CMs which are compliant with this specification.

C.2.1.1 Upstream Packet Classification Encoding

This field defines the parameters associated with an upstream Classifier.

Note that the same subtype fields defined are valid for both the encapsulated upstream and downstream packet classification configuration setting string. These type fields are not valid in other encoding contexts.

Type	Length	Value
22	n	

C.2.1.2 Downstream Packet Classification Encoding

This field defines the parameters associated with a downstream Classifier.

Note that the same subtype fields defined are valid for both the encapsulated upstream and downstream flow classification configuration setting string. These type fields are not valid in other encoding contexts.

Type	Length	Value
23	n	

C.2.1.3 General Packet Classifier Encodings

C.2.1.3.1 Classifier Reference

The value of the field specifies a reference for the Classifier. This value is unique per Dynamic Service message, configuration file, or Registration Request message.

Type	Length	Value
[22/23].1	1	

C.2.1.3.2 Classifier Identifier

The value of the field specifies an identifier for the Classifier. This value is unique to per Service Flow. The CMTS assigns the Packet Classifier Identifier.

Type	Length	Value
[22/23].2	2	

C.2.1.3.3 Service Flow Reference

The value of the field specifies a Service Flow Reference that identifies the corresponding Service Flow.

In all Packet Classifier TLVs that occur in any message where the Service Flow ID is not known (e.g. CM-initiated DSA-REQ and REG-REQ) this TLV **MUST** be included. In all Packet Classifier TLVs that occur in a DSC-REQ and CMTS-initiated DSA-REQ messages the Service Flow Reference **MUST NOT** be specified.

Type	Length	Value
[22/23].3	2	

C.2.1.3.4 Service Flow Identifier

The value of this field specifies the Service Flow ID that identifies the corresponding Service Flow.

In Packet Classifier TLVs where the Service Flow ID is not known, and this TLV **MUST NOT** be included (e.g. CM-initiated DSA-REQ and REG-REQ). In Packet Classifier TLVs that occur in a DSC-REQ and CMTS-initiated DSA-REQ message, the Service Flow ID **MUST** be specified.

Type	Length	Value
[22/23].4	4	

C.2.1.3.5 Rule Priority

The value of the field specifies the priority for the Classifier, which is used for determining the order of the Classifier. A higher value indicates higher priority.

Classifiers that appear in Configuration files and Registration messages **MAY** have priorities in the range 0 - 255 with the default value 0. Classifiers that appear in DSA/DSC message **MUST** have priorities in the range 64-191, with the default value 64.

Type	Length	Value
[22/23].5	1	

C.2.1.3.6 Classifier Activation State

The value of this field specifies whether this classifier should become active in selecting packets for the Service Flow. An inactive Classifier is typically used with an AdmittedQoSParameterSet to ensure resources are available for later activation

Type	Length	Value
[22/23].6	1	0 — Inactive 1 — Active

The default value is 1 — activate the classifier.

C.2.1.3.7 Dynamic Service Change Action

When received in a Dynamic Service Change Request, this indicates the action that should be taken with this classifier.

Type	Length	Value
[22/23].7	1	0 — DSC Add Classifier 1 — DSC Replace Classifier 2 — DSC Delete Classifier

C.2.1.4 Classifier Error Encodings

This field defines the parameters associated with Classifier Errors.

Type	Length	Value
[22/23].8	n	

A Classifier Error Parameter Set is defined by the following individual parameters: Errored Parameter, Confirmation Code and Error Message.

The Classifier Error Parameter Set is returned in REG-RSP, DSA-RSP and DSC-RSP messages to indicate the recipient's response to a Classifier establishment request in a REG-REQ, DSA-REQ or DSC-REQ message.

On failure, the sender **MUST** include one Classifier Error Parameter Set for each failed Classifier requested in the REG-REQ, DSA-REQ or DSC-REQ message. Classifier Error Parameter Set for the failed Classifier **MUST** include the Confirmation Code and Errored Parameter and **MAY** include an Error Message. If some Classifiers are successfully established, but others fail, Classifier Error Parameter Sets **MUST** only be sent for the failed Classifiers. On success of the entire transaction, the RSP or ACK message **MUST NOT** include a Classifier Error Parameter Set.

Multiple Classifier Error Parameter Sets **MAY** appear in a REG-RSP, DSA-RSP or DSC-RSP message, since multiple Classifier parameters may be in error. A message with even a single Classifier Error Parameter Set **MUST NOT** contain any other protocol Classifier Encodings (e.g. IP, 802.1P/Q).¹

A Classifier Error Parameter Set **MUST NOT** appear in any REG-REQ, DSA-REQ or DSC-REQ messages.

C.2.1.4.1 Errored Parameter

The value of this parameter identifies the subtype of a requested Classifier parameter in error in a rejected Classifier request or Service Class Name expansion response. A Classifier Error Parameter Set **MUST** have exactly one Errored Parameter TLV within a given Classifier Encoding.

Subtype	Length	Value
[22/23].8.1	n	Classifier Encoding Subtype in Error

If the length is one, then the value is the single-level subtype where the error was found, e.g. 7 indicates an invalid Change Action. If the length is two, then the value is the multi-level subtype where the error was found e.g. 9-2 indicates an invalid IP Protocol value.²

1. penultimate paragraph edited 06/22/99 per rfi-n-99043. ew

2. edited 06/22/99 per rfi-n-99043. ew

C.2.1.4.2 Error Code

This parameter indicates the status of the request. A non-zero value corresponds to the Confirmation Code as described in C.4. A Classifier Error Parameter Set **MUST** have exactly one Error Code within a given Classifier Encoding.

Subtype	Length	Value
[22/23].8.2	1	Confirmation code

A value of okay(0) indicates that the Classifier request was successful. Since a Classifier Error Parameter Set is only applies to errored parameters, this value **MUST NOT** be used.

C.2.1.4.3 Error Message

This subtype is optional in a Classifier Error Parameter Set. If present, it indicates a text string to be displayed on the CM console and/or log that further describes a rejected Classifier request. A Classifier Error Parameter Set **MAY** have zero or one Error Message subtypes within a given Classifier Encoding.

SubType	Length	Value
[22/23].8.3	n	Zero-terminated string of ASCII characters.

Note: The length N includes the terminating zero.

Note: The entire Classifier Encoding message must have a total length of less than 256 characters.

C.2.1.5 IP Packet Classification Encodings

This field defines the parameters associated with IP packet classification.

Type	Length	Value
[22/23].9	n	

C.2.1.5.1 IP Type of Service Range and Mask

The values of the field specify the matching parameters for the IP ToS byte range and mask. An IP packet with IP ToS byte value "ip-tos" matches this parameter if $\text{tos-low} \leq (\text{ip-tos AND tos-mask}) \leq \text{tos-high}$. If this field is omitted, then comparison of the IP packet ToS byte for this entry is irrelevant.

Type	Length	Value
[22/23].9.1	3	tos-low, tos-high, tos-mask

C.2.1.5.2 IP Protocol

The value of the field specifies the matching value for the IP Protocol field [RFC-1700]. If this parameter is omitted, then comparison of the IP header Protocol field for this entry is irrelevant.

There are two special IP Protocol field values: "256" matches traffic with any IP Protocol value, and "257" matches both TCP and UDP traffic. An entry that includes an IP Protocol field value greater than 257 **MUST** be invalidated for comparisons (i.e. no traffic can match this entry).

Type	Length	Value
[22/23].9.2	2	prot1, prot2

Valid Range
0 — 257

C.2.1.5.3 IP Source Address

The value of the field specifies the matching value for the IP source address. An IP packet with IP source address "ip-src" matches this parameter if $\text{src} = (\text{ip-src AND smask})$, where "smask" is the parameter from C.2.1.5.4. If this parameter is omitted, then comparison of the IP packet source address for this entry is irrelevant.

Type	Length	Value
[22/23].9.3	4	src1, src2, src3, src4

C.2.1.5.4 IP Source Mask

The value of the field specifies the mask value for the IP source address, as described in C.2.1.5.3. If this parameter is omitted, then the default IP source mask is 255.255.255.255.

Type	Length	Value
[22/23].9.4	4	smask1, smask2, smask3, smask4

C.2.1.5.5 IP Destination Address

The value of the field specifies the matching value for the IP destination address. An IP packet with IP destination address "ip-dst" matches this parameter if $\text{dst} = (\text{ip-dst AND dmask})$, where "dmask" is the parameter from C.2.1.5.6. If this parameter is omitted, then comparison of the IP packet destination address for this entry is irrelevant.

Type	Length	Value
[22/23].9.5	4	dst1, dst2, dst3, dst4

C.2.1.5.6 IP Destination Mask

The value of the field specifies the mask value for the IP destination address, as described in IP Destination Address. If this parameter is omitted, then the default IP destination mask is 255.255.255.255.

Type	Length	Value
[22/23].9.6	4	dmask1, dmask2, dmask3, dmask4

C.2.1.5.7 TCP/UDP Source Port Start

The value of the field specifies the low-end TCP/UDP source port value. An IP packet with TCP/UDP port value "src-port" matches this parameter if sportlow \leq src-port \leq sporthigh. If this parameter is omitted, then the default value of sportlow is 0. This parameter is irrelevant for non-TCP/UDP IP traffic.

Type	Length	Value
[22/23].9.7	2	sportlow1, sportlow2

C.2.1.5.8 TCP/UDP Source Port End

The value of the field specifies the high-end TCP/UDP source port value. An IP packet with TCP/UDP port value "src-port" matches this parameter if sportlow \leq src-port \leq sporthigh. If this parameter is omitted, then the default value of sporthigh is 65535. This parameter is irrelevant for non-TCP/UDP IP traffic.

Type	Length	Value
[22/23].9.8	2	sporthigh1, sporthigh2

C.2.1.5.9 TCP/UDP Destination Port Start

The value of the field specifies the low-end TCP/UDP destination port value. An IP packet with TCP/UDP port value "dst-port" matches this parameter if dportlow \leq dst-port \leq dporthigh. If this parameter is omitted, then the default value of dportlow is 0. This parameter is irrelevant for non-TCP/UDP IP traffic.

Type	Length	Value
[22/23].9.9	2	dportlow1, dportlow2

C.2.1.5.10 TCP/UDP Destination Port End

The value of the field specifies the high-end TCP/UDP destination port value. An IP packet with TCP/UDP port value "dst-port" matches this parameter if dportlow \leq dst-port \leq dporthigh. If this parameter is omitted, then the default value of dporthigh is 65535. This parameter is irrelevant for non-TCP/UDP IP traffic.

Type	Length	Value
[22/23].9.10	2	dporthigh1, dporthigh2

C.2.1.6 Ethernet LLC Packet Classification Encodings

This field defines the parameters associated with Ethernet LLC packet classification.

Type	Length	Value
[22/23].10	n	

C.2.1.6.1 Destination MAC Address

The values of the field specifies the matching parameters for the MAC destination address. An Ethernet packet with MAC destination address "etherdst" matches this parameter if dst = (etherdst AND msk). If this parameter is omitted, then comparison of the Ethernet MAC destination address for this entry is irrelevant.

Type	Length	Value
[22/23].10.1	12	dst1, dst2, dst3, dst4, dst5, dst6, msk1, msk2, msk3, msk4, msk5, msk6

C.2.1.6.2 Source MAC Address

The value of the field specifies the matching value for the MAC source address. If this parameter is omitted, then comparison of the Ethernet MAC source address for this entry is irrelevant.

Type	Length	Value
[22/23].10.2	6	src1, src2, src3, src4, src5, src6

C.2.1.6.3 Ethertype/IEEE 802.2 SAP

The values of the field specifies the matching parameters for the Ethernet Ethertype field (see [RFC1700]) or the 802.2 DSAP value. An DIX Ethernet packet with Ethertype "etype" (or an IEEE 802.3 packet with SNAP protocol "etype") matches this parameter if type = 1 and etype = eprot. An IEEE 802.3 packet with DSAP sap matches this parameter if type = 2 and sap = (eprot AND 255). If this parameter is omitted, then comparison of the Ethertype or 802.2 DSAP for this entry is irrelevant.

Type	Length	Value
[22/23].10.3	3	type, eprot1, eprot2

C.2.1.7 IEEE 802.1P/Q Packet Classification Encodings

This field defines the parameters associated with IEEE 802.1P/Q packet classification.

Type	Length	Value
[22/23].11	n	

C.2.1.7.1 IEEE 802.1P User_Priority

The values of the field specify the matching parameters for the IEEE 802.1P user_priority bits. An Ethernet packet with IEEE 802.1P user_priority value "priority" matches these parameters if pri-low <= priority <= pri-high. If this field is omitted, then comparison of the IEEE 802.1P user_priority bits for this entry is irrelevant.

If this parameter is specified for an entry, then Ethernet packets without IEEE 802.1Q encapsulation MUST NOT match this entry. If this parameter is specified for an entry on a CM that does not support forwarding of IEEE 802.1Q encapsulated traffic, then this entry MUST NOT be used for any traffic.

Type	Length	Value
[22/23].11.1	2	pri-low, pri-high

Valid Range
0 — 7 for pri-low and pri-high

C.2.1.7.2 IEEE 802.1Q VLAN_ID

The value of the field specify the matching value for the IEEE 802.1Q vlan_id bits. Only the first (i.e. left-most) 12 bits of the specified vlan_id field are significant; the final four bits must be ignored for comparison. If this field is omitted, then comparison of the IEEE 802.1Q vlan_id bits for this entry is irrelevant.

If this parameter is specified for an entry, then Ethernet packets without IEEE 802.1Q encapsulation **MUST NOT** match this entry. If this parameter is specified for an entry on a CM that does not support forwarding of IEEE 802.1Q encapsulated traffic, then this entry **MUST NOT** be used for any traffic.

Type	Length	Value
[22/23].11.2	2	vlan_id1, vlan_id2

C.2.1.7.3 Vendor Specific Classifier Parameters

This allows vendors to encode vendor-specific Classifier parameters. The Vendor ID **MUST** be the first TLV embedded inside Vendor Specific Classifier Parameters. If the first TLV inside Vendor Specific Classifier Parameters is not a Vendor ID, then the TLV must be discarded. (Refer to C.1.1.17)¹

Type	Length	Value
[22/23].43	n	

C.2.1.8 Upstream-Specific Classification Encodings**C.2.1.8.1 Classifier Activation Signal**

This field **MUST** only be used in Dynamic Service Change messages that originate from the CMTS and which affect the Active parameter set. It is not present in any other Service Flow signaling messages.

Type	Length	Value
22.12	1	1 — Activate/Deactivate Classifier on Request 2 — Activate/Deactivate Classifier on Ack

This field directs the modem to change its upstream transmission characteristics to match those in the DSC either immediately on receiving the DSC-Request, or only after receiving the DSC-Ack. In particular, it signals the time of (de-)activation of any classifiers which are changed by this DSC exchange.

The default value is 2 for a bandwidth increase. The default value is 1 for a bandwidth decrease. If increase or decrease is ambiguous, then the default value is 2.

1. edited 06/22/99 per rfi-n-99043. ew

C.2.2 Service Flow Encodings

The following type/length/value encodings **MUST** be used in the configuration file, registration messages, and Dynamic Service messages to encode parameters for Service Flows. All multi-octet quantities are in network-byte order, i.e., the octet containing the most-significant bits is the first transmitted on the wire.

The following configuration settings **MUST** be supported by all CMs which are compliant with this specification.

C.2.2.1 Upstream Service Flow Encodings

This field defines the parameters associated with upstream scheduling for a Service Flow. It is somewhat complex in that it is composed from a number of encapsulated type/length/value fields.

Note that the encapsulated upstream and downstream Service Flow configuration setting strings share the same subtype field numbering plan, because many of the subtype fields defined are valid for both types of configuration settings. These type fields are not valid in other encoding contexts.

Type	Length	Value
24	n	

C.2.2.2 Downstream Service Flow Encodings

This field defines the parameters associated with downstream scheduling for a Service Flow. It is somewhat complex in that it is composed from a number of encapsulated type/length/value fields.

Note that the encapsulated upstream and downstream flow classification configuration setting strings share the same subtype field numbering plan, because many of the subtype fields defined are valid for both types of configuration settings except Service Flow encodings.

Type	Length	Value
25	n	

C.2.2.3 General Service Flow Encodings

C.2.2.3.1 Service Flow Reference

The Service Flow Reference is used to associate a packet classifier encoding with a Service Flow encoding. A Service Flow Reference is only used to establish a Service Flow ID. Once the Service Flow exists and has an assigned Service Flow ID, the Service Flow Reference **MUST** no longer be used.

Type	Length	Value
[24/25].1	2	

C.2.2.3.2 Service Flow Identifier

The Service Flow Identifier is used by the CMTS as the primary reference of a Service Flow. Only the CMTS can issue a Service Flow Identifier. It uses this parameterization to issue Service Flow Identifiers in CMTS-initiated DSA/DSC-Requests and in its REG/DSA/DSC-Response to CM-initiated REG/DSA/DSC-Requests. The CM specifies the SFID of a service flow using this parameter in a DSC-REQ message.

The configuration file **MUST NOT** contain this parameter.

Type	Length	Value
[24/25].2	4	

C.2.2.3.3 Service Identifier

The value of this field specifies the Service Identifier assigned by the CMTS to a Service Flow with a non-null AdmittedQosParameterSet or ActiveQosParameterSet. This is used in the bandwidth allocation MAP to assign upstream bandwidth. This field **MUST** be present in CMTS-initiated DSA-REQ or DSC-REQ message related to establishing an admitted or active upstream Service Flow. This field **MUST** also be present in REG-RSP, DSA-RSP and DSC-RSP messages related to the successful establishment of an admitted or active upstream Service Flow.

Even though a Service Flow has been successfully admitted or activated (i.e. has an assigned Service ID) the Service Flow ID **MUST** be used for subsequent DSx message signalling as it is the primary handle for a service flow. If a Service Flow is no longer admitted or active (via DSC-REQ) its Service ID **MAY** be reassigned by the CMTS.

SubType	Length	Value
[24/25].3	2	SID (low-order 14 bits)

C.2.2.3.4 Service Class Name

The value of the field refers to a predefined CMTS service configuration to be used for this Service Flow.

Type	Length	Value
[24/25].4	n	Zero-terminated string of ASCII characters.

Note: The length 'n' includes the terminating zero.

When the Service Class Name is used in a Service Flow encoding, it indicates that all the unspecified QoS Parameters of the Service Flow need to be provided by the CMTS. It is up to the operator to synchronize the definition of Service Class Names in the CMTS and in the configuration file.

C.2.2.4 Service Flow Error Encodings

This field defines the parameters associated with Service Flow Errors.

Type	Length	Value
[24/25].5	n	

A Service Flow Error Parameter Set is defined by the following individual parameters: Confirmation Code, Errored Parameter and Error Message.

The Service Flow Error Parameter Set is returned in REG-RSP, DSA-RSP and DSC-RSP messages to indicate the recipient's response to a Service Flow establishment request in a REG-REQ, DSA-REQ or DSC-REQ message. The Service Flow Error Parameter Set is returned in REG-ACK, DSA-ACK and DSC-ACK messages to indicate the recipient's response to the expansion of a Service Class Name in a corresponding REG-RSP, DSA-RSP or DSC-RSP.

On failure, the sender **MUST** include one Service Flow Error Parameter Set for each failed Service Flow requested in the REG-REQ, DSA-REQ or DSC-REQ message. On failure, the sender **MUST** include one Service Flow Error Parameter Set for each failed Service Class Name expansion in the REG-RSP, DSA-RSP or DSC-RSP message. Service Flow Error Parameter Set for the failed Service Flow **MUST** include the Confirmation Code and Errored Parameter and **MAY** include an Error Message. If some Service Flows are successfully established, but others fail, Service Flow Error Parameter Sets are only **REQUIRED** for the failed service flows, but **MAY** be included for successful Service Flows.

On success of the entire transaction, the RSP or ACK message **MUST NOT** include a Service Flow Error Parameter Set.

Multiple Service Flow Error Parameter Sets **MAY** appear in a REG-RSP, DSA-RSP, DSC-RSP, REG-ACK, DSA-ACK or DSC-ACK message, since multiple Service Flow parameters may be in error. A message with even a single Service Flow Error Parameter Set **MUST NOT** contain any QoS Parameters.

A Service Flow Error Parameter Set **MUST NOT** appear in any REG-REQ, DSA-REQ or DSC-REQ messages.

C.2.2.4.1 Errored Parameter

The value of this parameter identifies the subtype of a requested Service Flow parameter in error in a rejected Service Flow request or Service Class Name expansion response. A Service Flow Error Parameter Set **MUST** have exactly one Errored Parameter TLV within a given Service Flow Encoding.

Subtype	Length	Value
[24/25].5.1	1	Service Flow Encoding Subtype in Error

C.2.2.4.2 Error Code

This parameter indicates the status of the request. A non-zero value corresponds to the Confirmation Code as described in C.4. A Service Flow Error Parameter Set **MUST** have exactly one Error Code within a given Service Flow Encoding.

Subtype	Length	Value
[24/25].5.2	1	Confirmation code

A value of okay(0) indicates that the Service Flow request was successful. Since a Service Flow Error Parameter Set is only applies to errored parameters, this value **MUST NOT** be used.

C.2.2.4.3 Error Message

This subtype is optional in a Service Flow Error Parameter Set. If present, it indicates a text string to be displayed on the CM console and/or log that further describes a rejected Service Flow request. A Service Flow Error Parameter Set MAY have zero or one Error Message subtypes within a given Service Flow Encoding.

SubType	Length	Value
[24/25].5.3	n	Zero-terminated string of ASCII characters.

Note: The length N includes the terminating zero.

Note: The entire Service Flow Encoding message must have a total length of less than 256 characters.

C.2.2.5 Common Upstream and Downstream Quality-of-Service Parameter Encodings

The remaining Type 24 & 25 parameters are QoS Parameters. Any given QoS Parameter type MUST appear zero or one times per Service Flow Encoding.

C.2.2.5.1 Quality of Service Parameter Set Type

This parameter MUST appear within every Service Flow Encoding set. It specifies the proper application of the QoS Parameter Set: to the Provisioned set, the Admitted set, and/or the Active set. Various combinations may be specified with this parameter.

Type	Length	Value
[24/25].6	1	Bit # 0 Provisioned Set
		Bit # 1 Admitted Set
		Bit # 2 Active Set
		Bit Values
		000 Reserved
		001 Apply to Provisioned set only
		010 Perform admission control and apply to admitted set
		011 Apply to Provisioned and admitted set, and perform admission control
		100 Check against admitted set, perform admission control if needed, activate, and apply to active set combinations
		101 Apply to provisioned and active sets, perform admission control, and activate this Service flow ¹
		110 Perform admission control and activate, apply parameters to both admitted and active sets.
		111 Apply to Provisioned, admitted, and active sets; perform admission control and activate this Service Flow.

A Service Flow Encoding that appears in a Registration-Request message MUST specify a ProvisionedQoSParameterSet, and MAY also specify an Admitted and/or Active set. A Service Flow Encoding that appears in a Dynamic Service message MUST NOT specify the ProvisionedQoSParameterSet.

A single DSA/DSC message MAY contain multiple Service Flow Configuration Setting TLVs for the same Service Flow. A CMTS MUST handle a single update to up to each of the three parameter sets. The ability to process multiple Configuration Setting TLVs that specify the same QoS parameter set is NOT required, and is

1. edited 06/22/99 per rfi-n-99043 ew

left as a vendor-specific function. If a DSA/DSC contains multiple updates to a single QoS parameter set and the vendor does not support such updates, then the CMTS MUST reply with error code 2-reject-unrecognized-configuration-setting.

C.2.2.5.2 Traffic Priority

The value of this parameter specifies the priority assigned to a Service Flow. Given two Service Flows identical in all QoS parameters besides priority, the higher priority Service Flow SHOULD be given lower delay and higher buffering preference. For otherwise non-identical Service Flows, the priority parameter SHOULD NOT take precedence over any conflicting Service Flow QoS parameter. The specific algorithm for enforcing this parameter is not mandated here.

For upstream service flows, the CMTS SHOULD use this parameter when determining precedence in request service and grant generation, and the CM MUST preferentially select contention Request opportunities for Priority Request Service IDs (refer to A.2.3) based on this priority and its Request/Transmission Policy (refer to C.2.2.6.3).

Type	Length	Value
[24/25].7	1	0 to 7 — Higher numbers indicate higher priority

Note: The default priority is 0.

C.2.2.5.3 Maximum Sustained Traffic Rate

This parameter is the rate parameter R of a token-bucket-based rate limit for packets. R is expressed in bits per second, and must take into account all MAC frame data PDU of the Service Flow from the byte following the MAC header HCS to the end of the CRC¹. The number of bytes forwarded-(in bytes) is limited during any time interval T by Max(T), as described in the expression

$$\text{Max}(T) = T * (R / 8) + B, \quad (1)$$

where the parameter B (in bytes) is the Maximum Traffic Burst Configuration Setting (refer to C.2.2.5.4).

Note: This parameter does not limit the instantaneous rate of the Service Flow.

Note: The specific algorithm for enforcing this parameter is not mandated here. Any implementation which satisfies the above equation is conformant.

Note: If this parameter is omitted or set to zero, then there is no explicitly-enforced traffic rate maximum. This field specifies only a bound, not a guarantee that this rate is available.

C.2.2.5.3.1 Upstream Maximum Sustained Traffic Rate

For an upstream Service Flow, the CM MUST NOT request bandwidth exceeding the Max(T) requirement in (1) during any interval T because this could force the CMTS to fill MAPs with deferred grants.

The CM MUST defer upstream packets that violate (1) and “rate shape” them to meet the expression, up to a limit as implemented by vendor buffering restrictions.

1. The payload size includes every PDU in a Concatenated MAC Frame.

The CMTS **MUST** enforce expression (1) on all upstream data transmissions, including data sent in contention. The CMTS **MAY** consider unused grants in calculations involving this parameter. The CMTS **MAY** enforce this limit by any of the following methods: (a) discarding over-limit requests, (b) deferring (through zero-length grants) the grant until it is conforming to the allowed limit, or (c) discarding over-limit data packets. A CMTS **MUST** report this condition to a policy module. If the CMTS is policing by discarding either packets or requests, the CMTS **MUST** allow a margin of error between the CM and CMTS algorithms.

Type	Length	Value
24.8	4	R (in bits per second)

C.2.2.5.3.2 Downstream Maximum Sustained Traffic Rate

For a downstream Service Flow, this parameter is only applicable at the CMTS. The CMTS **MUST** enforce expression (1) on all downstream data transmissions. The CMTS **MUST NOT** forward downstream packets that violates (1) in any interval T. The CMTS **SHOULD** "rate shape" the downstream traffic by enqueueing packets arriving in excess of (1), and delay them until the expression can be met.

This parameter is not intended for enforcement on the CM.

Type	Length	Value
25.8	4	R (in bits per second)

C.2.2.5.4 Maximum Traffic Burst

The value of this parameter specifies the token bucket size B (in bytes) for this Service Flow as described in expression (1). This value is calculated from the byte following the MAC header HCS to the end of the CRC¹.

If this parameter is omitted, then the default B is 1522 bytes. The minimum value of B is the larger of 1522 bytes or the value of Maximum Concatenated Burst Size (refer to C.2.2.6.1).

Type	Length	Value
[24/25].9	4	B (bytes)

Note: The specific algorithm for enforcing this parameter is not mandated here. Any implementation which satisfies the above equation is conformant.

C.2.2.5.5 Minimum Reserved Traffic Rate

This parameter specifies the guaranteed minimum rate, in bits/sec, reserved for this Service Flow. This value is calculated from the byte following the MAC header HCS to the end of the CRC². If this parameter is omitted, then the default minimum rate is 0 bits/sec (i.e., no bandwidth is reserved for the flow by default).

This field is only applicable at the CMTS and **MUST** be enforced by the CMTS.

Type	Length	Value
[24/25].10	4	

Note: The specific algorithm for enforcing the value specified in this field is not mandated here.

1. The payload size includes every PDU in a Concatenated MAC Frame.
2. The payload size includes every PDU in a Concatenated MAC Frame.

C.2.2.5.6 Assumed Minimum Reserved Rate Packet Size

The value of this field specifies an assumed minimum packet size (in bytes) for which the Minimum Reserved Traffic Rate will be provided. This parameter is defined in bytes and is specified as the bytes following the MAC header HCS to the end of the CRC¹. If the Service Flow sends packets of size smaller than this specified value, such packets will be treated as being of the size specified in this parameter for calculating the Minimum Reserved Traffic Rate.

The CMTS MUST apply this parameter to its Minimum Reserved Traffic Rate algorithm. This parameter is used by the CMTS to estimate the per packet overhead of each packet in the service flow.

If this parameter is omitted, then the default value is CMTS implementation dependent.²

Type	Length	Value
[24/25].11	2	

C.2.2.5.7 Timeout for Active QoS Parameters

The value of this parameter specifies the maximum duration resources remain unused on an active Service Flow. If there is no activity on the Service Flow within this time interval, the CMTS MUST change the active and admitted QoS Parameter Sets to null. The CMTS MUST signal this resource change with a DSC-REQ to the CM.

If defined, this parameter MUST be enforced at the CMTS and SHOULD NOT be enforced at the CM.

Type	Length	Value
[24/25].12	2	seconds

The value of 0 means that the flow is of infinite duration and MUST NOT be timed out due to inactivity. The default value is 0.

C.2.2.5.8 Timeout for Admitted QoS Parameters

The value of this parameter specifies the duration that the CMTS MUST hold resources for a Service Flow's Admitted QoS Parameter Set while they are in excess of its Active QoS Parameter Set. If there is no DSC-REQ to activate the Admitted QoS Parameter Set within this time interval, the resources that are admitted but not activated MUST be released, and only the active resources retained. The CMTS MUST set the Admitted QoS Parameter Set equal to the Active QoS Parameter Set for the Service Flow and initiate a DSC-REQ exchange with the CM to inform it of the change.

If this parameter is omitted, then the default value is 200 seconds. The value of 0 means that the Service Flow can remain in the admitted state for an infinite amount of time and MUST NOT be timed out due to inactivity. However, this is subject to policy control by the CMTS.

This parameter MUST be enforced by the CMTS. The CMTS MAY set the response value less than the requested value.

Type	Length	Value
[24/25].13	2	seconds

1. The payload size includes every PDU in a Concatenated MAC Frame.

2. sentence edited 06/22/99 per rfi-n-99043. ew

C.2.2.5.9 Vendor Specific QoS Parameters

This allows vendors to encode vendor-specific QoS parameters. The Vendor ID **MUST** be the first TLV embedded inside Vendor Specific QoS Parameters. If the first TLV inside Vendor Specific QoS Parameters is not a Vendor ID, then the TLV must be discarded. (Refer to C.1.1.17)¹

Type	Length	Value
[24/25].43	n	

C.2.2.6 Upstream-Specific QoS Parameter Encodings**C.2.2.6.1 Maximum Concatenated Burst**

The value of this parameter specifies the maximum concatenated burst (in bytes) which a Service Flow is allowed. This parameter is calculated from the FC byte of the Concatenation MAC Header to the last CRC in the concatenated MAC frame.

A value of 0 means there is no limit. The default value is 0.

This field is only applicable at the CM. If defined, this parameter **MUST** be enforced at the CM.

Note: This value does not include any physical layer overhead.

Type	Length	Value
24.14	2	

Note: This applies only to concatenated bursts. It is legal and, in fact, it may be useful to set this smaller than the maximum Ethernet packet size. Of course, it is also legal to set this equal to or larger than the maximum Ethernet packet size.

C.2.2.6.2 Service Flow Scheduling Type

The value of this parameter specifies which upstream scheduling service is used for upstream transmission requests and packet transmissions. If this parameter is omitted, then the Best Effort service **MUST** be assumed.

This parameter is only applicable at the CMTS. If defined, this parameter **MUST** be enforced by the CMTS.

Type	Length	Value
24.15	1	0 Reserved
		1 for Undefined (CMTS implementation-dependent ²)
		2 for Best Effort
		3 for Non-Real-Time Polling Service
		4 for Real-Time Polling Service
		5 for Unsolicited Grant Service with Activity Detection
		6 for Unsolicited Grant Service
		7 through 255 are reserved for future use

1. edited 06/22/99 per rfi-n-99043 ew

2. The specific implementation dependent scheduling service type could be defined in the 24.43 Vendor Specific Information Field.

C.2.2.6.3 Request/Transmission Policy

The value of this parameter specifies which IUC opportunities the CM uses for upstream transmission requests and packet transmissions for this Service Flow. Requests for this Service Flow may be piggybacked with data and whether data packets transmitted on this Service Flow can be concatenated, fragmented or payload header suppressed. If this parameter is omitted, then the CM MUST use all possible transmission opportunities presented in Section 7.1.2, MUST assume that piggybacking requests is allowed, and MUST use the concatenation, fragmentation and payload header suppression policies established through the modem capabilities exchange.

This field is only applicable at the CM. If defined, this parameter MUST be enforced at the CM.

Type	Length	Value
24.16	4	Bit #0 The Service Flow MUST NOT use "all CMs" broadcast request opportunities. Bit #1 The Service Flow MUST NOT use Priority Request multicast request opportunities. (Refer to A.2.3) Bit #2 The Service Flow MUST NOT use Request/Data opportunities for Requests Bit #3 The Service Flow MUST NOT use Request/Data opportunities for Data Bit #4 The Service Flow MUST NOT piggyback requests with data. Bit #5 The Service Flow MUST NOT concatenate data. Bit #6 The Service Flow MUST NOT fragment data Bit #7 The Service Flow MUST NOT suppress payload headers Bit #8 ¹ The Service Flow MUST drop packets that do not fit in the Unsolicited Grant Size ² All other bits are reserved.

Note: Data grants include both short and long data grants.

C.2.2.6.4 Nominal Polling Interval

The value of this parameter specifies the nominal interval (in units of microseconds) between successive unicast request opportunities for this Service Flow on the upstream channel. This parameter is typically suited for Real-Time and Non-Real-Time Polling Service.

The ideal schedule for enforcing this parameter is defined by a reference time t_0 , with the desired transmission times $t_i = t_0 + i \cdot \text{interval}$. The actual poll times, t'_i MUST be in the range $t_i \leq t'_i \leq t_i + \text{jitter}$, where interval is the value specified with this TLV, and jitter is Tolerated Poll Jitter. The accuracy of the ideal poll times, t_i , are measured relative to the CMTS Master Clock used to generate timestamps (refer to Section 7.3).

This field is only applicable at the CMTS. If defined, this parameter MUST be enforced by the CMTS.

Type	Length	Value
24.17	4	μsec

1. This bit only applies to Service Flows with the Unsolicited Grant Service Flow Scheduling Type, if this bit is set on any other Service Flow Scheduling type it MUST be ignored
2. Packets that classify to an Unsolicited Grant Service Flow and are larger than the Grant Size associated with that Service Flow are normally transmitted on the Primary Service Flow. This parameter overrides that default behavior.

C.2.2.6.5 Tolerated Poll Jitter

The values in this parameter specifies the maximum amount of time that the unicast request interval MAY be delayed from the nominal periodic schedule (measured in microseconds) for this Service Flow.

The ideal schedule for enforcing this parameter is defined by a reference time t_0 , with the desired poll times $t_i = t_0 + i \cdot \text{interval}$. The actual poll, t'_i MUST be in the range $t_i \leq t'_i \leq t_i + \text{jitter}$, where jitter is the value specified with this TLV and interval is the Nominal Poll Interval. The accuracy of the ideal poll times, t_i , are measured relative to the CMTS Master Clock used to generate timestamps (refer to Section 7.3).

This parameter is only applicable at the CMTS. If defined, this parameter represents a service commitment (or admission criteria) at the CMTS.

Type	Length	Value
24.18	4	μsec

C.2.2.6.6 Unsolicited Grant Size

The value of this parameter specifies the unsolicited grant size in bytes. The grant size includes the entire MAC frame data PDU from the Frame Control byte to end of the MAC frame.

This parameter is applicable at the CMTS and MUST be enforced at the CMTS.

Type	Length	Value
24.19	2	

Note: For UGS, this parameter should be used by the CMTS to compute the size of the unsolicited grant in minislots.

C.2.2.6.7 Nominal Grant Interval

The value of this parameter specifies the nominal interval (in units of microseconds) between successive data grant opportunities for this Service Flow. This parameter is required for Unsolicited Grant and Unsolicited Grant with Activity Detection Service Flows. This parameter is optional for Real-Time Polling Service Flows. If the ActiveQoSParameterSet is Null, however, the CMTS MUST initiate a DSD-REQ exchange with the CM to inform it of the deleted Service Flow.

The ideal schedule for enforcing this parameter is defined by a reference time t_0 , with the desired transmission times $t_i = t_0 + i \cdot \text{interval}$. The actual grant times, t'_i MUST be in the range $t_i \leq t'_i \leq t_i + \text{jitter}$, where interval is the value specified with this TLV, and jitter is the Tolerated Grant Jitter. When an upstream Service Flow with either Unsolicited Grant or Unsolicited Grant with Activity Detection scheduling becomes active, the first grant MUST define the start of this interval, i.e. the first grant MUST be for an ideal transmission time, t_i . When multiple grants per interval are requested, all grants MUST be within this interval, thus the Nominal Grant Interval and Tolerated Grant Jitter MUST be maintained by the CMTS for all grants in this Service Flow. The accuracy of the ideal grant times, t_i , are measured relative to the CMTS Master Clock used to generate timestamps (refer to Section 7.3).

This field is mandatory for Unsolicited Grant and Unsolicited Grant with Activity Detection Scheduling Types. This field is only applicable at the CMTS, and MUST be enforced by the CMTS.

Type	Length	Value
24.20	4	μsec

C.2.2.6.8 Tolerated Grant Jitter

The values in this parameter specifies the maximum amount of time that the transmission opportunities MAY be delayed from the nominal periodic schedule (measured in microseconds) for this Service Flow.

The ideal schedule for enforcing this parameter is defined by a reference time t_0 , with the desired transmission times $t_i = t_0 + i \cdot \text{interval}$. The actual transmission opportunities, t'_i MUST be in the range $t_i \leq t'_i \leq t_i + \text{jitter}$, where jitter is the value specified with this TLV and interval is the Nominal Grant Interval. The accuracy of the ideal grant times, t_i , are measured relative to the CMTS Master Clock used to generate timestamps (refer to Section 7.3).

This field is mandatory for Unsolicited Grant and Unsolicited Grant with Activity Detection Scheduling Types. This field is only applicable at the CMTS, and MUST be enforced by the CMTS.

Type	Length	Value
24.21	4	μsec

C.2.2.6.9 Grants per Interval

The value of this parameter indicates the number of data grants per Nominal Grant Interval. This is intended to enable the addition of sessions to an existing Unsolicited Grant Service Flow via the Dynamic Service Change mechanism, without negatively impacting existing sessions.

The ideal schedule for enforcing this parameter is defined by a reference time t_0 , with the desired transmission times $t_i = t_0 + i \cdot \text{interval}$. The actual grant times, t'_i MUST be in the range $t_i \leq t'_i \leq t_i + \text{jitter}$, where interval is the Nominal Grant Interval, and jitter is the Tolerated Grant Jitter. When multiple grants per interval are requested, all grants MUST be within this interval, thus the Nominal Grant Interval and Tolerated Grant Jitter MUST be maintained by the CMTS for all grants in this Service Flow.

This field is mandatory for Unsolicited Grant and Unsolicited Grant with Activity Detection Scheduling Types. This field is only applicable at the CMTS, and MUST be enforced by the CMTS.

Type	Length	Value	Valid Range
24.22	1	# of grants	0-127 ¹

C.2.2.6.10 IP Type of Service Overwrite

The CMTS MUST overwrite IP packets with IP ToS byte value "orig-ip-tos" with the value "new-ip-tos", where $\text{new-ip-tos} = ((\text{orig-ip-tos} \text{ AND } \text{tos-and-mask}) \text{ OR } \text{tos-or-mask})$. If this parameter is omitted, then the IP packet ToS byte is not overwritten.

This parameter is only applicable at the CMTS. If defined, this parameter MUST be enforced by the CMTS.

Type	Length	Value
24.23	2	tos-and-mask, tos-or-mask

(section C.2.2.6.11 deleted 06/22/99 per rfi-n-99042. ew)

1. text added 06/22/99 per rfi-n-99043

C.2.2.7 Downstream-Specific QoS Parameter Encodings

C.2.2.7.1 Maximum Downstream Latency

The value of this parameter specifies the maximum latency between the reception of a packet by the CMTS on its NSI and the forwarding of the packet to its RF Interface.

If defined, this parameter represents a service commitment (or admission criteria) at the CMTS and **MUST** be guaranteed by the CMTS. A CMTS does not have to meet this service commitment for Service Flows that exceed their minimum downstream reserved rate.

Type	Length	Value
25.14	4	μsec^1

C.2.2.8 Payload Header Suppression

This field defines the parameters associated with Payload Header Suppression.

Type	Length	Value
26	n	

Note: The entire Payload Header Suppression TLV must have a length of less than 255 characters.

C.2.2.8.1 Classifier Reference

The value of the field specifies a Classifier Reference that identifies the corresponding Classifier. (Refer to C.2.1.3.1)

Type	Length	Value
26.1	1	

C.2.2.8.2 Classifier Identifier

The value of the field specifies a Classifier Identifier that identifies the corresponding Classifier. (Refer to C.2.1.3.2)

Type	Length	Value
26.2	2	

C.2.2.8.3 Service Flow Reference

The value of the field specifies a Service Flow Reference that identifies the corresponding Service Flow. (Refer to C.2.2.3.1)

Type	Length	Value
26.3	2 ²	

1. edited 06/22/99 per rfi-n-99043 ew

2. edited 06/22/99 per rfi-n-99043 ew

C.2.2.8.4 Service Flow Identifier

The value of the field specifies a Service Flow Identifier that identifies the corresponding Service Flow. All downstream PHS Rules **MUST** use the Service Flow Identifier of the Primary Downstream Service Flow. (Refer to C.2.2.3.2)

Type	Length	Value
26.4	4	

C.2.2.8.5 Dynamic Service Change Action

When received in a Dynamic Service Change Request, this indicates the action that **MUST** be taken with this payload header suppression byte string.

Type	Length	Value
26.5	1	0 — Add PHS Rule 1 — Set PHS Rule 2 — Delete PHS Rule 3 — Delete all PHS Rules

The “Set PHS Rule” command is used to add the specific TLV’s for an undefined payload header suppression rule. It **MUST NOT** be used to modify existing TLV’s.

Note: When deleting all PHS Rules any corresponding Payload Header Suppression Index **MUST** be ignored.

Note: An attempt to Add a PHS Rule which already exists is an error condition.

C.2.2.9 Payload Header Suppression Error Encodings

This field defines the parameters associated with Payload Header Suppression Errors.

Type	Length	Value
26.6	n	

A Payload Header Suppression Error Parameter Set is defined by the following individual parameters: Errored Parameter, Confirmation Code and Error Message.

The Payload Header Suppression Error Parameter Set is returned in REG-RSP, DSA-RSP and DSC-RSP messages to indicate the recipient’s response to a Payload Header Suppression Rule establishment request in a REG-REQ, DSA-REQ or DSC-REQ message.

On failure, the sender **MUST** include one Payload Header Suppression Error Parameter Set for each failed Payload Header Suppression Rule requested in the REG-REQ, DSA-REQ or DSC-REQ message. Payload Header Suppression Error Parameter Set for the failed Payload Header Suppression Rule **MUST** include the Confirmation Code and Errored Parameter and **MAY** include an Error Message. If some Payload Header Suppression Rules are successfully established, but others fail, Payload Header Suppression Error Parameter Sets **MUST** only be sent for the failed Payload Header Suppression Rules. On success of the entire transaction, the RSP or ACK message **MUST NOT** include a Payload Header Suppression Error Parameter Set.

Multiple Payload Header Suppression Error Parameter Sets **MAY** appear in a REG-RSP, DSA-RSP or DSC-RSP message, since multiple Payload Header Suppression parameters may be in error. A message with even a single Payload Header Suppression Error Parameter Set **MUST NOT** contain any other protocol Payload Header Suppression Encodings (e.g. IP, 802.1P/Q).¹

A Payload Header Suppression Error Parameter Set **MUST NOT** appear in any REG-REQ, DSA-REQ or DSC-REQ messages.

C.2.2.9.1 Errored Parameter

The value of this parameter identifies the subtype of a requested Payload Header Suppression parameter in error in a rejected Payload Header Suppression request or Service Class Name expansion response. A Payload Header Suppression Error Parameter Set **MUST** have exactly one Errored Parameter TLV within a given Payload Header Suppression Encoding.

Subtype	Length	Value
26.6.1	1	Payload Header Suppression Encoding Subtype in Error

C.2.2.9.2 Error Code

This parameter indicates the status of the request. A non-zero value corresponds to the Confirmation Code as described in C.4. A Payload Header Suppression Error Parameter Set **MUST** have exactly one Error Code within a given Payload Header Suppression Encoding.

Subtype	Length	Value
26.6.2	1	Confirmation code

A value of okay(0) indicates that the Payload Header Suppression request was successful. Since a Payload Header Suppression Error Parameter Set only applies to errored parameters, this value **MUST NOT** be used.

C.2.2.9.3 Error Message

This subtype is optional in a Payload Header Suppression Error Parameter Set. If present, it indicates a text string to be displayed on the CM console and/or log that further describes a rejected Payload Header Suppression request. A Payload Header Suppression Error Parameter Set **MAY** have zero or one Error Message subtypes within a given Payload Header Suppression Encoding.

SubType	Length	Value
26.6.3	n	Zero-terminated string of ASCII characters.

Note: The length n includes the terminating zero.

Note: The entire Payload Header Suppression Encoding message must have a total length of less than 256 characters.

C.2.2.10 Payload Header Suppression Rule Encodings

C.2.2.10.1 Payload Header Suppression Field (PHSF)

The value of this field are the bytes of the headers which **MUST** be suppressed by the sending entity, and **MUST** be restored by the receiving entity.

Type	Length	Value
26.7	n	string of bytes suppressed

1. edited 06/22/99 per rfi-n-99043. ew

The length *n* MUST always be the same as the value for PHSS.

C.2.2.10.2 Payload Header Suppression Index (PHSI)

The Payload Header Suppression Index (PHSI) has a value between 1 and 255 which uniquely references the suppressed byte string. The Index is unique per Service Flow in the upstream direction and unique per CM in the downstream direction. The upstream and downstream PHSI values are independent of each other.

Type	Length	Value
26.8	1	index value

C.2.2.10.3 Payload Header Suppression Mask (PHSM)

The value of this field is used to interpret the values in the Payload Header Suppression Field. It is used at both the sending and receiving entities on the link. The PHSM allows fields such as sequence numbers or checksums which vary in value to be excluded from suppression with the constant bytes around them suppressed.

Type	Length	Value
26.9	<i>n</i>	bit 0: 0 = don't suppress first byte; 1 = suppress first byte bit 1: 0 = don't suppress second byte; 1 = suppress second byte bit <i>x</i> : 0 = don't suppress (<i>x</i> +1) byte; 1 = suppress (<i>x</i> +1) byte

The length *n* is ceiling(PHSS/8). Bit 0 is the MSB of the Value field. The value of each sequential bit in the PHSM is an attribute for the corresponding sequential byte in the PHSF.

If the bit value is a "1", the sending entity should suppress the byte, and the receiving entity should restore the byte from its cached PHSF. If the bit value is a "0", the sending entity should not suppress the byte, and the receiving entity should restore the byte by using the next byte in the packet.

If this TLV is not included, the default is to suppress all bytes.

C.2.2.10.4 Payload Header Suppression Size (PHSS)

The value of this field is the total number of bytes in the header to be suppressed and then restored in a Service Flow that uses Payload Header Suppression.

Type	Length	Value
26.10	1	number of bytes in the suppression string

This TLV is used when a Service Flow is being created. All packets which get classified and assigned to a Service Flow with Payload Header Suppression enabled MUST suppress the specified number of bytes. If this TLV is not included in a Service Flow definition, or is included with a value of 0 bytes, then Payload Header Suppression is disabled. A non-zero value indicates Payload Header Suppression is enabled.

C.2.2.10.5 Payload Header Suppression Verification (PHSV)

The value of this field indicates to the sending entity whether it MUST verify the PHSF string in the packet with its cache entry before doing the substitution.

Type	Length	Value
26.11	1	0 = verify 1 = don't verify

If this TLV is not included, the default is to verify. Only the sender **MUST** verify suppressed bytes. If verification fails, the Payload Header **MUST NOT** be suppressed. (Refer to Section 8.4.3)

C.2.2.10.6 Vendor Specific PHS Parameters

This allows vendors to encode vendor-specific PHS parameters. The Vendor ID **MUST** be the first TLV embedded inside Vendor Specific PHS Parameters. If the first TLV inside Vendor Specific PHS Parameters is not a Vendor ID, then the TLV must be discarded. (Refer to C.1.1.17)¹

Type	Length	Value
26.43 ²	n	

1. edited 06/22/99 per rfi-n-99043.ew

2. edited 06/22/99 per rfi-n-99043.ew

C.2.3 Parameter Applicability for Upstream Service Scheduling

Service Flow Parameter	Unsolicited Grant	Real-Time Polling	Unsolicited Grant w/Activity Detection	Non-Real-Time Polling	Best Effort
Traffic Priority	N/A ^a	N/A	N/A	Optional Default is 0	Optional Default is 0
Maximum Sustained Traffic Rate	Optional ^b	Optional Default is 0	Optional Default is 0	Optional Default is 0	Optional Default is 0
Maximum Traffic Burst Size	N/A	Optional Default is 1522	Optional Default is 1522	Optional Default is 1522	Optional Default is 1522
Minimum Reserved Traffic Rate	N/A	Optional Default is 0	Optional Default is 0	Optional Default is 0	Optional Default is 0
Unsolicited Grant Size	Mandatory	Optional ^c	Mandatory	N/A	N/A
Assumed Min Reserved Rate Packet Size ^d	Optional Default is CMTS implementation dependent	Optional Default is CMTS implementation dependent	Optional Default is CMTS implementation dependent	Optional Default is CMTS implementation dependent	Optional Default is CMTS implementation dependent ^e
Maximum Concatenated Burst	N/A	Optional	Optional	Optional	Optional
Upstream Flow Scheduling Service Type	Mandatory	Mandatory	Mandatory	Mandatory	Optional Default is 1 (Best Effort)
Request/Transmission Policy	Optional Default is 127	Optional Default is 31	Optional Default is 127	Optional Default is 127	Optional Default is 0
Nominal Polling Interval	N/A	Mandatory	Optional Default is Nominal Data Grant Interval	Optional Default is CMTS-specific	N/A
Tolerated Poll Jitter	N/A	Optional Default is CMTS-specific	Optional Default is CMTS-specific	N/A	N/A
Nominal Grant Interval	Mandatory	Optional	Mandatory	N/A	N/A
Tolerated Grant Jitter	Optional Default is CMTS-specific	N/A	Optional Default is CMTS-specific	N/A	N/A
Grants per Interval	Mandatory	Optional	Mandatory	N/A	N/A

a.N/A means Not Applicable to this service flow scheduling type. If included in a request for a service flow of this service flow scheduling type this request **MUST** be denied.

b.This generally only makes sense for controlling piggybacked requests

c.If this parameter is not present then there are no unsolicited grants with this service

d.Edited 06/22/99 per rfi-n-99043. ew

e.Row entries edited 06/22/99 per rfi-n-99043. ew

C.3 Encodings for Other Interfaces

C.3.1 Telephone Settings Option

This configuration setting describes parameters which are specific to telephone return systems. It is composed from a number of encapsulated type/length/value fields. See [DOCSIS6].

Type	Length	Value
15 (= TRI_CFG01)	n	

C.3.2 Baseline Privacy Configuration Settings Option

This configuration setting describes parameters which are specific to Baseline Privacy. It is composed from a number of encapsulated type/length/value fields. See [DOCSIS8].

Type	Length	Value
17 (= BP_CFG)	n	

C.4 Confirmation Code

The Confirmation Code (CC) provides a common way to indicate failures for Registration Response, Registration Ack, Dynamic Service Addition-Response, Dynamic Service Addition-Ack, Dynamic Service Delete-Response, Dynamic Service Change-Response and Dynamic Service Change-Ack MAC Management Messages.

Confirmation Code is one of the following:

- okay / success(0)
- reject-other(1)
- reject-unrecognized-configuration-setting(2)
- reject-temporary / reject-resource(3)
- reject-permanent / reject-admin(4)
- reject-not-owner(5)
- reject-service-flow-not-found(6)
- reject-service-flow-exists(7)
- reject-required-parameter-not-present(8)
- reject-header-suppression(9)
- reject-unknown-transaction-id(10)
- reject-authentication-failure(11)

The Confirmation Codes MUST be used in the following way:

- Okay or success(0) means the message was received and successful.
- Reject-other(1) is used when none of the other reason codes apply.
- Reject-unrecognized-configuration setting(2) is used when a configuration setting is not recognized or when its value is outside of the specified range.
- Reject-temporary(3), also known as reject-resource, indicates that the current loading of the CMTS or CM prevents granting the request, but that the request might succeed at another time.
- Reject-permanent(4), also known as reject-admin, indicates that, for policy, configuration, or capabilities reasons, the request would never be granted unless the CMTS or CM were manually reconfigured or replaced.
- Reject-not-owner(5) the requester is not associated with this service flow.
- Reject-service-flow-not-found(6) the Service Flow indicated in the request does not exist.
- Reject-service-flow-exists(7) the Service Flow to be added already exists.
- Reject-required-parameter-not-present(8) a required parameter has been omitted.
- Reject-header-suppression(9) the requested header suppression cannot be supported for whatever reason.
- Reject-unknown-transaction-id(10) the requested transaction continuation is invalid because the receiving end-point does not view the transaction as being 'in process' (i.e. the message is unexpected or out of order).
- Reject-authentication-failure(11) the requested transaction was rejected by the authorization module.¹

1. bullet added 06/22/99per rfi-n-99043. ew

This page intentionally left blank.

Appendix D. CM Configuration Interface Specification

D.1 CM IP Addressing

D.1.1 DHCP Fields Used by the CM

The following fields **MUST** be present in the DHCP request from the CM and **MUST** be set as described below:

- The hardware type (htype) **MUST** be set to 1 (Ethernet).
- The hardware length (hlen) **MUST** be set to 6.
- The client hardware address (chaddr) **MUST** be set to the 48 bit MAC address associated with the RF interface of the CM.
- The “client identifier” option **MUST** be included, with the hardware type set to 1, and the value set to the same 48 bit MAC address as the chaddr field.
- The “parameter request list” option **MUST** be included. The option codes that **MUST** be included in the list are:
 - Option code 1 (Subnet Mask)
 - Option code 2 (Time Offset)
 - Option code 3 (Router Option)
 - Option code 4 (Time Server Option)
 - Option code 7 (Log Server Option)
 - Option code 60 (Vendor Specific Option) — to allow for the differentiation between DOCSIS 1.1 and DOCSIS 1.0 CM requests, a compliant CM **MUST** send the following ASCII coded string in Option code 60, “docsis1.1:xxxxxxx”. Where xxxxx **MUST** be the hexadecimal encoding of the Modem Capabilities, refer to C.1.3.1.

The following fields are expected in the DHCP response returned to the CM. The CM **MUST** configure itself based on the DHCP response.

- The IP address to be used by the CM (yiaddr).
- The IP address of the TFTP server for use in the next phase of the bootstrap process (siaddr).
- If the DHCP server is on a different network (requiring a relay agent), then the IP address of the relay agent (giaddr). Note: this may differ from the IP address of the first hop router.
- The name of the CM configuration file to be read from the TFTP server by the CM (file).
- The subnet mask to be used by the CM (Subnet Mask, option 1).
- The time offset of the CM from Universal Coordinated Time (UTC) (Time Offset, option 2). This is used by the CM to calculate the local time for use in time-stamping error logs.
- A list of addresses of one or more routers to be used for forwarding CM-originated IP traffic (Router Option, option 3). The CM is not required to use more than one router IP address for forwarding.
- A list of [RFC-868] time-servers from which the current time may be obtained (Time Server Option, option 4).
- A list of SYSLOG servers to which logging information may be sent (Log Server Option, option 7); see [DOCSIS5].

D.2 CM Configuration

D.2.1 CM Binary Configuration File Format

The CM-specific configuration data **MUST** be contained in a file which is downloaded to the CM via TFTP. This is a binary file in the same format defined for DHCP vendor extension data [RFC-2132].

It **MUST** consist of a number of configuration settings (1 per parameter) each of the form

Type Length Value

Where Type is a single-octet identifier which defines the parameter

Length is a single octet containing the length of the value field in octets (not including type and length fields)

Value is from one to 254 octets containing the specific value for the parameter

The configuration settings **MUST** follow each other directly in the file, which is a stream of octets (no record markers).

Configuration settings are divided into three types:

- Standard configuration settings which **MUST** be present
- Standard configuration settings which **MAY** be present
- Vendor-specific configuration settings.

CMs **MUST** be capable of processing all standard configuration settings. CMs **MUST** ignore any configuration setting present in the configuration file which it cannot interpret. To allow uniform management of CM's conformant to this specification, conformant CM's **MUST** support a 8192-byte configuration file at a minimum.

Authentication of the provisioning information is provided by two message integrity check (MIC) configuration settings, CM MIC and CMTS MIC.

- CM MIC is a digest which ensures that the data sent from the provisioning server were not modified en route. This is **NOT** an authenticated digest (it does not include any shared secret).
- CMTS MIC is a digest used to authenticate the provisioning server to the CMTS during registration. It is taken over a number of fields one of which is a shared secret between the CMTS and the provisioning server.

Use of the CM MIC allows the CMTS to authenticate the provisioning data without needing to receive the entire file.

Thus the file structure is of the form shown in Figure D-1:




Figure D-1. Binary Configuration File Format

D.2.2 Configuration File Settings

The following configuration settings **MUST** be included in the configuration file and **MUST** be supported by all CMs.

- Network Access Configuration Setting
- CM MIC Configuration Setting
- CMTS MIC Configuration Setting
- End Configuration Setting
- DOCSIS 1.0 Class of Service Configuration Setting¹

— or —

- Upstream Service Flow Configuration Setting
- Downstream Service Flow Configuration Setting

The following configuration settings **MAY** be included in the configuration file and if present **MUST** be supported by all CMs.

- Downstream Frequency Configuration Setting
- Upstream Channel ID Configuration Setting
- Baseline Privacy Configuration Setting
- Software Upgrade Filename Configuration Setting
- Upstream Packet Classification Setting
- Downstream Packet Classification Setting
- SNMP Write-Access Control
- SNMP MIB Object
- Software Server IP Address
- CPE Ethernet MAC Address
- Maximum Number of CPEs
- Maximum Number of Classifiers
- Privacy Enable Configuration Setting
- Payload Header Suppression
- TFTP Server Timestamp
- TFTP Server Provisioned Modem Address
- Pad Configuration Setting

The following configurations **MAY** be included in the configuration file and if present, and applicable to this type of modem, **MUST** be supported.

- Telephone Settings Option

1. A DOCSIS 1.0 CM **MUST** be provided with a DOCSIS 1.0 Class of Service Configuration. A CM conformant with this specification should only be provisioned with DOCSIS 1.0 Class of Service Configuration information if it is to behave as a DOCSIS 1.0 CM, otherwise it **MUST** be provisioned with Service Flow Configuration Settings.

The following configuration settings MAY be included in the configuration file and if present MAY be supported by a CM.

- Vendor-Specific Configuration Settings

Note: There is a limit on the size of registration request and registration response frames (see section 6.2.5.2). The configuration file MUST NOT cause the CM to CMTS to exceed that limit.¹

D.2.3 Configuration File Creation

The sequence of operations required to create the configuration file is as shown in Figure D-2 through Figure D-5.

1. Create the type/length/value entries for all the parameters required by the CM.

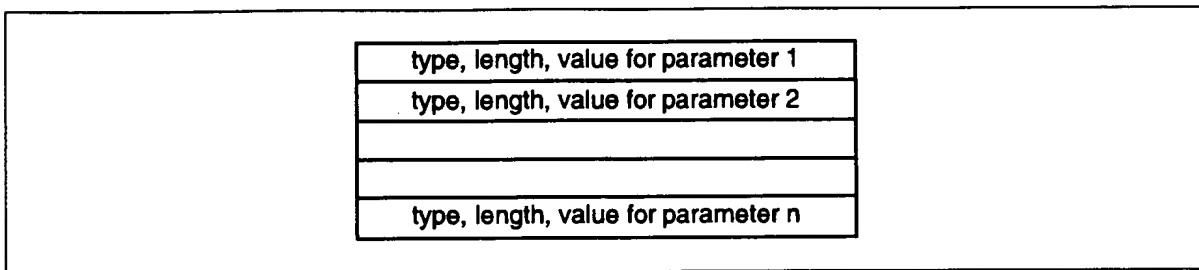


Figure D-2. Create TLV Entries for Parameters Required by the CM

2. Calculate the CM message integrity check (MIC) configuration setting as defined in Section D.2.3.1 and add to the file following the last parameter using code and length values defined for this field.

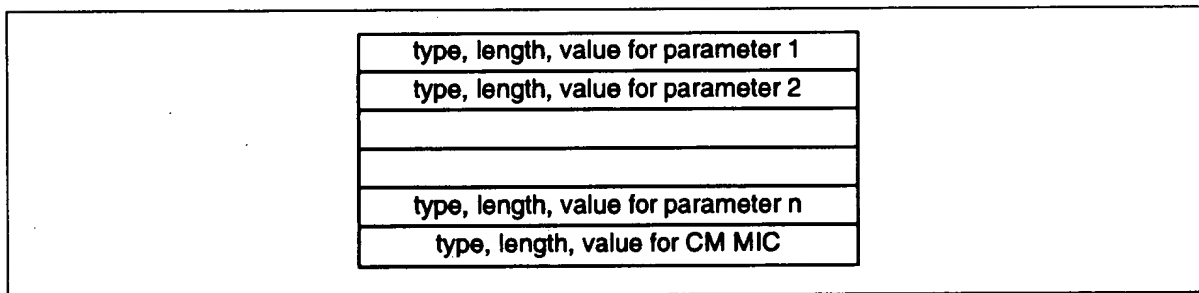


Figure D-3. Add CM MIC

3. Calculate the CMTS message integrity check (MIC) configuration setting as defined in Section D.3.1 and add to the file following the CM MIC using code and length values defined for this field.

1. Final paragraph, Section D.2.2. added 6/7/99 per ECN rfi-n-99035

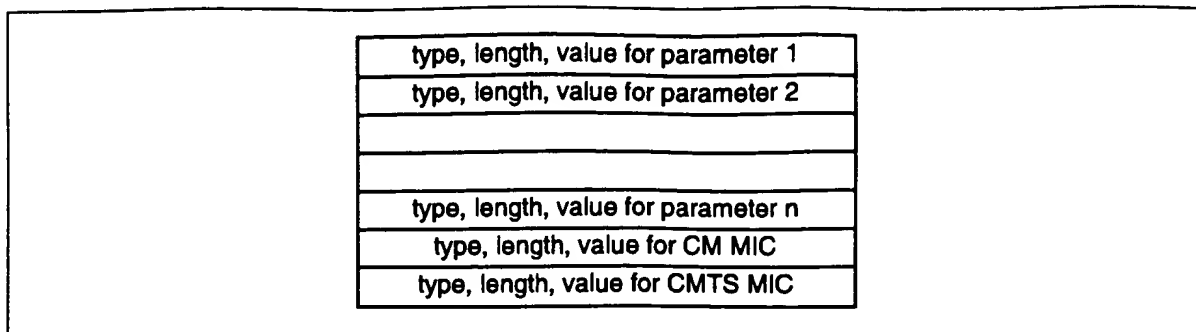


Figure D-4. Add CMTS MIC

4. Add the end of data marker.

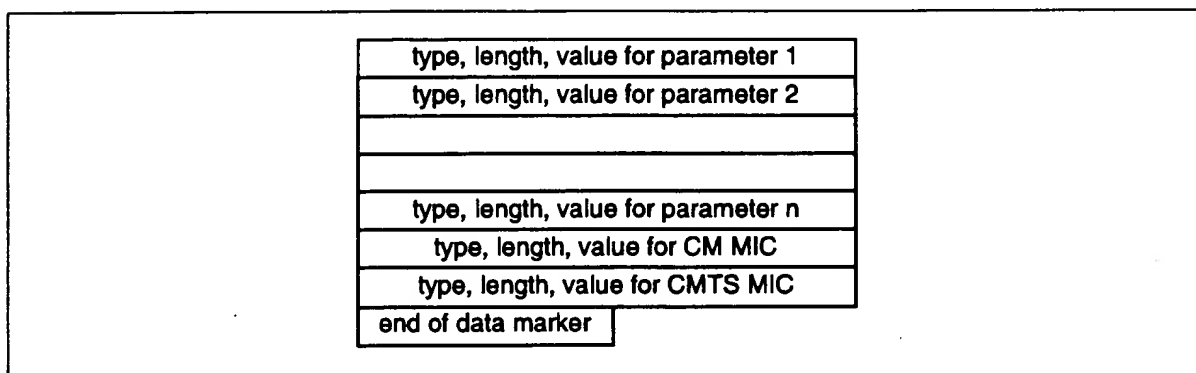


Figure D-5. Add End of Data Marker

D.2.3.1 CM MIC Calculation

The CM message integrity check configuration setting **MUST** be calculated by performing an MD5 digest over the bytes of the configuration setting fields. It is calculated over the bytes of these settings as they appear in the TFTPed image, without regard to TLV ordering or contents. There are two exceptions to this disregard of the contents of the TFTPed image:

1. The bytes of the CM MIC TLV itself are omitted from the calculation. This includes the type, length, and value fields.
2. The bytes of the CMTS MIC TLV are omitted from the calculation. This includes the type, length, and value fields.

On receipt of a configuration file, the CM **MUST** recompute the digest and compare it to the CM MIC configuration setting in the file. If the digests do not match then the configuration file **MUST** be discarded

D.3 Configuration Verification

It is necessary to verify that the CM's configuration file has come from a trusted source. Thus, the CMTS and the configuration server share an Authentication String that they use to verify portions of the CM's configuration in the Registration Request.

D.3.1 CMTS MIC Calculation

The CMTS message integrity check configuration setting **MUST** be calculated by performing an MD5 digest over the following configuration setting fields, when present in the configuration file, in the order shown:

- Downstream Frequency Configuration Setting
- Upstream Channel ID Configuration Setting
- Network Access Configuration Setting
- DOCSIS 1.0 Class of Service Configuration Setting
- Baseline Privacy Configuration Setting
- Vendor-Specific Configuration Settings
- CM MIC Configuration Setting
- Maximum Number of CPEs
- TFTP Server Timestamp
- TFTP Server Provisioned Modem Address
- Upstream Packet Classification Setting
- Downstream Packet Classification Setting
- Upstream Service Flow Configuration Setting
- Downstream Service Flow Configuration Setting
- Maximum Number of Classifiers
- Privacy Enable Configuration Setting
- Payload Header Suppression

The bulleted list specifies the order of operations when calculating the CMTS MIC over configuration setting Type fields. The CMTS **MUST** calculate the CMTS MIC over TLVs of the same Type in the order they were received. Within Type fields, the CMTS **MUST** calculate the CMTS MIC over the Subtypes in the order they were received. To allow for correct CMTS MIC calculation by the CMTS, the CM **MUST NOT** reorder configuration file TLVs of the same Type or Subtypes within any given Type in its Registration-Request message.

All configuration setting fields **MUST** be treated as if they were contiguous data when calculating the CM MIC.¹

The digest **MUST** be added to the configuration file as its own configuration setting field using the CMTS MIC Configuration Setting encoding.

The authentication string is a shared secret between the provisioning server (which creates the configuration files) and the CMTS. It allows the CMTS to authenticate the CM provisioning. The authentication string is to be used as the key for calculating the keyed CMTS MIC digest as stated in D.3.1.1.²

The mechanism by which the shared secret is managed is up to the system operator.

On receipt of a configuration file, the CM **MUST** forward the CMTS MIC as part of the registration request (REG-REQ).

1. edited per rfi-n-99053 06/21/99

2. edited per rfi-n-99043 06/21/99

On receipt of a REG-REQ, the CMTS **MUST** recompute the digest over the included fields and the authentication string and compare it to the CMTS MIC configuration setting in the file. If the digests do not match, the registration request **MUST** be rejected by setting the authentication failure result in the registration response status field.

D.3.1.1 Digest Calculation

The CMTS MIC digest field **MUST** be calculated using HMAC-MD5 as defined in [RFC-2104].

This page intentionally left blank.

Appendix E. MAC Service Definition

This section is informational. In case of conflict between this section and any normative section of this specification, the normative section takes precedence.

E.1 MAC Service Overview

The DOCSIS MAC provides a protocol service interface to upper-layer services. Examples of upper-layer services include a DOCSIS bridge, embedded applications (e.g. Packetcable/VOIP), a host interface (e.g. NIC adapter with NDIS driver), and layer three routers (e.g. IP router).

The MAC Service interface defines the functional layering between the upper layer service and the MAC. As such it defines the functionality of the MAC which is provided by the underlying MAC protocols. This interface is a protocol interface, not a specific implementation interface.

The following data services are provided by the MAC service interface:

- A MAC service exists for classifying and transmitting packets to MAC service flows.
- A MAC service exists for receiving packets from MAC service flows. Packets may be received with suppressed headers.
- A MAC service exists for transmitting and receiving packets with suppressed headers. The headers of transmitted packets are suppressed based upon matching classifier rules. The headers of received suppressed packets are regenerated based upon a packet header index negotiated between the CM and CMTS.
- A MAC service exists for synchronization of grant timing between the MAC and the upper layer service. This clock synchronization is required for applications such as embedded Packetcable VOIP clients in which the packetization period needs to be synchronized with the arrival of scheduled grants from the CMTS.
- A MAC service exists for synchronization of the upper layer clock with the CMTS Controlled Master Clock.

It should be noted that a firewall and policy based filtering service may be inserted between the MAC layer and the upper layer service, but such a service is not modeled in this MAC service definition.

The following control services are provided by the MAC service interface:

- A MAC service exists for the upper layer to learn of the existence of provisioned service flows and QoS traffic parameter settings at registration time.
- A MAC service exists for the upper layer to create service flows. Using this service the upper layer initiates the admitted/activated QoS parameter sets, classifier rules, and packet suppression headers for the service flow.
- A MAC service exists for the upper layer to delete service flows.
- A MAC service exists for the upper layer to change service flows. Using this service the upper layer modifies the admitted/activated QoS parameter sets, classifier rules, and packet suppression headers.
- A MAC service exists for controlling the classification of and transmission of PDUs with suppressed headers. At most a single suppressed header is defined for a single classification rule. The upper layer service is responsible for defining both the definition of suppressed headers (including wild-card don't-suppress fields) and the unique classification rule that discriminates each header. In addition to the classification rule, the MAC service can perform a full match of all remaining header bytes to prevent generation of false headers if so configured by the upper layer service.
- A MAC service exists for controlling two-phase control of QoS traffic resources. Two phase activation is controlled by the upper layer service provide both admitted QoS parameters and active QoS parameters within the appropriate service request. Upon receipt of an affirmative indication the upper layer service

knows that the admitted QoS parameter set has been reserved by the CMTS, and that the activated QoS parameter set has been activated by the CMTS. Barring catastrophic failure (such as resizing of the bandwidth of the upstream PHY), admitted resources will be guaranteed to be available for activation, and active resources will be guaranteed to be available for use in packet transmission.

A control function for locating an unused service flow and binding it or a specific identified service flow to a specific upper layer service may also exist. The details of such a function are not specified and are implementation dependent.

Other control functions may exist at the MAC service interface, such as functions for querying the status of active service flows and packet classification tables, or functions from the MAC service to the upper layer service to enable the upper layer service to authorize service flows requested by the peer MAC layer service, but those functions are not modeled in this MAC service definition.

Other MAC services that are not service flow related also exist, such as functions for controlling the MAC service MAC address and SAID multicast filtering functions, but those functions are not modeled in this MAC service definition.

E.1.1 MAC Service Parameters

The MAC service utilizes the following parameters. For a full description of the parameters consult the Theory of Operation and other relevant sections within the body of the RFI specification.

- Service Flow QoS Traffic Parameters

MAC activate-service-flow and change-service-flow primitives allow common, upstream, and downstream QoS traffic parameters to be provided. When such parameters are provided they override whatever values were configured for those parameters at provisioning time or at the time the service flow was created by the upper layer service.

- Active/Reserved QoS Traffic Parameters

If two-phase service flow activation is being used, then two complete sets of QoS Traffic Parameters are controlled. The admitted QoS Parameters state the requirements for reservation of resources to be authorized by the CMTS. The activated QoS Parameters state the requirements for activation of resources to be authorized by the CMTS. Admitted QoS parameters may be activated at a future time by the upper layer service. Activated QoS parameters may be used immediately by the upper layer service.

- Service Flow Classification Filter Rules

Zero or more classification filter rules may be provided for each service flow that is controlled by the upper layer service. Classifiers are identified with a classifier identifier.

- Service Flow PHS Suppressed Headers

Zero or more PHS suppressed header strings with their associated verification control and mask variables may be defined for each service flow. When such headers are defined, they are associated 1-to-1 with specific classification rules. In order to regenerate packets with suppressed headers a payload header suppression index is negotiated between the CM and CMTS.

E.2 MAC Data Service Interface

MAC services are defined for transmission and reception of data to and from service flows. Typically an upper layer service will utilize service flows for mapping of various classes of traffic to different service flows. Mappings to service flows may be defined for low priority traffic, high priority traffic, and multiple special traffic classes such as constant bit rate traffic which is scheduled by periodic grants from the CMTS at the MAC layer.

The following specific data service interfaces are provided by the MAC service to the upper layer service. These represent an abstraction of the service provided and do not imply a particular implementation:

```
MAC_DATA.request
MAC_DATA.indicate
MAC_GRANT_SYNCHRONIZE.indicate
MAC_CMTS_MASTER_CLOCK_SYNCHRONIZE.indicate
```

E.2.1 MAC_DATA.request

Issued by the upper-layer service to request classification and transmission of an IEEE 802.3 or DIX formatted PDU to the RF.

Parameters:

- PDU - IEEE 802.3 or DIX encoded PDU including all layer two header fields and optional FCS. PDU is the only mandatory parameter.
- padding - is used when the PDU is less than 60 bytes and it is desired to maintain ISO8802-3 transparency.
- ServiceFlowID - if included the MAC service circumvents the packet classification function and maps the packet to the specific service flow indicated by the ServiceFlowID value.
- ServiceClassName, RulePriority - if included this tuple identifies the service class name of an active service flow to which the packet is to be mapped so long as a classifier does not exist at a rule priority higher than the rule priority supplied.

Expanded Service Description:

Transmit a PDU from upper-layer service to MAC/PHY. The only mandatory parameter is PDU. PDU contains all layer-2 headers, layer-3 headers, data, and (optional) layer-2 checksum.

If PDU is the only parameter, the packet is subjected to the MAC packet classification filtering function in order to determine how the packet is mapped to a specific service flow. The results of the packet classification operation determine on which service flow the packet is to be transmitted and whether or not the packet should be transmitted with suppressed headers.

If the parameter ServiceFlowID is supplied the packet can be directed to the specifically identified service flow.

If the parameter tuple ServiceClassName, RulePriority is supplied the packet is directed to the first active service flow that matches the service class name so long as a classifier does not exist at a rule priority higher than the rule priority supplied. This service is used by upper layer policy enforcers to allow zero or more dynamic rules to be matched for selected traffic (e.g. voice) while all other traffic is forced to a service flow within the named ServiceFlowClass. If no active service flow with the Service Class Name exists, then the service perform normal packet classification.

In all cases, if no classifier match is found, or if none of the combinations of parameters maps to a specific service flow, the packet will be directed to the primary service flow.

The following pseudo code describes the intended operation of the MAC_DATA.request service interface:

MAC_DATA.request

PDU

[ServiceFlowID]

[ServiceClassName, RulePriority]

FIND_FIRST_SERVICE_FLOW_ID (ServiceClassName) returns ServiceFlowID of first service flow whose ServiceClassName equals the parameter of the procedure or NULL if no matching service flow found.

SEARCH_CLASSIFIER_TABLE (PriorityRange) searches all rules within the specified priority range and returns either the ServiceFlowID associated with the rule or NULL if no classifier rule found.

TxServiceFlowID = NULL

IF (ServiceFlowID DEFINED)

 TxServiceFlowID = MAC_DATA.ServiceFlowID

ELSEIF (ServiceClassName DEFINED and RulePriority DEFINED)

 TxServiceFlowID = FIND_FIRST_SERVICE_FLOW_ID (ServiceClassName)

 SearchID = SEARCH_CLASSIFIER_TABLE (All Priority Levels)

 IF (SearchID not NULL and ClassifierRule.Priority >= MAC_DATA.RulePriority)

 TxServiceFlowID = SearchID

ELSE [PDU only]

 TxServiceFlow = SEARCH_CLASSIFIER_TABLE (All Priority Levels)

IF (TxServiceFlowID = NULL)

 TRANSMIT_PDU (PrimaryServiceFlowID)

ELSE

 TRANSMIT_PDU (TxServiceFlowID)

E.2.2 MAC_DATA.Indicate

Issued by the MAC to indicate reception of an IEEE 802.3 or DIX PDU for the upper-layer service from the RF.

Parameters:

- PDU - IEEE 802.3 or DIX encoded PDU including all layer two header fields and FCS.

E.2.3 MAC_GRANT_SYNCHRONIZE.Indicate

Issued by the MAC service to the upper layer service to indicate the timing of grant arrivals from the CTMS. It is not stated how the upper layer derives the latency if any between the reception of the indication and the actual arrival of grants (within the bounds of permitted grant jitter) from the CMTS. It should be noted that in UGS applications it is expected that the MAC layer service will increase the grant rate or decrease the grant rate based upon the number of grants per interval QoS traffic parameter. It should also be noted that as the number of grants per interval is increased or decreased that the timing of grant arrivals will change also. It should also be noted that when synchronization is achieved with the CMTS downstream master clock, this indication may only be required once per active service flow. No implication is given as to how this function is implemented.

Parameters:

- ServiceFlowID - unique identifier value for the specific active service flow receiving grants.

E.2.4 MAC_CMTS_MASTER_CLOCK_SYNCHRONIZE.indicate

Issued by the MAC service to the upper layer service to indicate the timing of the CMTS master clock. No implication is given as to how often or how many times this indication is delivered by the MAC service to the upper layer service. No implication is given as to how this function is implemented.

Parameters:

- No parameters specified.

E.3 MAC Control Service Interface

A collection of MAC services are defined for control of MAC service flows and classifiers. It should be noted that an upper layer service may use these services to provide an upper layer traffic construct such as "connections" or "subflows" or "micro-flows". However, except for the ability to modify individual classifiers, no explicit semantics is defined for such upper layer models. Thus control of MAC service flow QoS parameters is specified in the aggregate.

The following specific control service interface functions are provided by the MAC service to the upper layer service. These represent an abstraction of the service provided and do not imply a particular implementation:

MAC_REGISTRATION_RESPONSE.indicate
MAC_CREATE_SERVICE_FLOW.request/response/indicate
MAC_DELETE_SERVICE_FLOW.request/response/indicate
MAC_CHANGE_SERVICE_FLOW.request/response/indicate

E.3.1 MAC_REGISTRATION_RESPONSE.indicate¹

Issued by the DOSCIS MAC to the upper layer service to indicate the complete set service flows and service flow QoS traffic parameters that have been provisioned and authorized by the registration phase of the MAC. Subsequent changes to service flow activation state or addition and deletion of service flows are communicated to the upper layer service with indications from the other MAC control services.

Parameters:

- Registration TLVs - any and all TLVs that are needed for service flow and service flow parameter definition including provisioned QoS parameters. See the normative body of the specification for more details.

E.3.2 MAC_CREATE_SERVICE_FLOW.request

Issued by the upper-layer service to the MAC to request the creation of a new service flow within the MAC service. This primitive is not issued for service flows that are configured and registered, but rather for dynamically created service flows. This primitive may also define classifiers for the service flow and supply admitted and activated QoS parameters. This function invokes DSA signaling.

Parameters:

- ServiceFlowID - unique id value for the specific service flow being created.
- ServiceClassName - service flow class name for the service flow being created.

1. Title changed per rfi-n-99043 06/21/99 ew

- Admitted QoS Parameters - zero or more upstream, downstream, and common traffic parameters for the service flow.
- Activated QoS Parameters - zero or more upstream, downstream, and common traffic parameters for the service flow.

E.3.3 MAC_CREATE_SERVICE_FLOW.response

Issued by the MAC service to the upper layer service to indicate the success or failure of the request to create a service flow.

Parameters:

- ServiceFlowID - unique identifier value for the specific service flow being created.
- ResponseCode - success or failure code

E.3.4 MAC_CREATE_SERVICE_FLOW.indicate

Issued by the MAC service to notify the upper-layer service of the creation of a new service flow within the MAC service. This primitive is not issued for service flows that have been administratively pre-configured, but rather for dynamically defined service flows. In this draft of the specification this notification is advisory only.

Parameters:

- ServiceFlowID - unique id value for the specific service flow being created.
- ServiceClassName - service flow class name for the service flow being created.
- Admitted QoS Parameters - zero or more upstream, downstream, and common traffic parameters for the service flow.
- Activated QoS Parameters - zero or more upstream, downstream, and common traffic parameters for the service flow.

E.3.5 MAC_DELETE_SERVICE_FLOW.request

Issued by the upper-layer service to the MAC to request the deletion of a service flow and all authorized and admitted QoS parameters within the MAC service. This function invokes DSD signaling.

Parameters:

- ServiceFlowID - optional unique identifier value for the deleted service flow.

E.3.6 MAC_DELETE_SERVICE_FLOW.response

Issued by the MAC service to the upper layer service to indicate the success or failure of the request to delete a service flow.

Parameters:

- ServiceFlowID - unique identifier value for the specific service flow being deleted.
- ResponseCode - success or failure code

E.3.7 MAC_DELETE_SERVICE_FLOW.Indicate

Issued by the MAC service to notify the upper-layer service of deletion of a service flow within the MAC service.

Parameters:

- ServiceFlowID - optional unique identifier value for the deleted service flow.

E.3.8 MAC_CHANGE_SERVICE_FLOW.request

Issued by the upper-layer service to the MAC to request modifications to a specific created and acquired service flow. This function is able to define both the complete set of classifiers and incremental changes to classifiers (add/remove). This function defines the complete set of admitted and active QoS parameters for a service flow. This function invokes DSC MAC-layer signaling.

Parameters:

- ServiceFlowID - unique identifier value for the specific service flow being modified.
- zero or more packet classification rules with add/remove semantics and LLC, IP, and 802.1pq parameters.
- Admitted QoS Parameters - zero or more upstream, downstream, and common traffic parameters for the service flow.
- Activated QoS Parameters - zero or more upstream, downstream, and common traffic parameters for the service flow.

E.3.9 MAC_CHANGE_SERVICE_FLOW.response

Issued by the MAC service to the upper layer service to indicate the success or failure of the request to change a service flow.

Parameters:

- ServiceFlowID - unique identifier value for the specific service flow being released.
- ResponseCode - success or failure code

E.3.10 MAC_CHANGE_SERVICE_FLOW.Indicate

Issued by the DOSCIS MAC service to notify upper-layer service of a request to change a service flow. In this specification the notification is advisory only and no confirmation is required before the service flow is changed. Change-service-flow indications are generated based upon DSC signaling. DSC signaling can be originated based upon change-service-flow events between the peer upper-layer service and its MAC service, or based upon network resource failures such as a resizing of the total available bandwidth at the PHY layer. How the upper layer service reacts to forced reductions in admitted or reserved QoS traffic parameters is not specified.

Parameters:

- ServiceFlowID - unique identifier for the service flow being activated.
- packet classification rules with LLC, IP, and 802.1pq parameters, and with zero or more PHS_CLASSIFIER_IDENTIFIERS.
- Admitted QoS Parameters - zero or more upstream, downstream, and common traffic parameters for the service flow.

- Activated QoS Parameters - zero or more upstream, downstream, and common traffic parameters for the service flow.

E.4 MAC Service Usage Scenarios

Upper layer entities utilize the services provided by the MAC in order to control service flows and in order to send and receive data packets. The partition of function between the upper-layer-service and the MAC service is demonstrated by the following scenarios.

E.4.1 Transmission of PDUs from Upper Layer Service to MAC DATA Service

- Upper layer service transmits PDUs via the MAC_DATA service.
- MAC_DATA service classifies transmitted PDUs using the classification table, and transmits the PDUs on the appropriate service flow. The classification function may also cause the packet header to be suppressed according to a header suppression template stored with the classification rule. It is possible for the upper layer service to circumvent this classification function.
- MAC_DATA service enforces all service flow based QoS traffic shaping parameters.
- MAC_DATA service transmits PDUs on DOCSIS RF as scheduled by the MAC layer.

E.4.2 Reception of PDUs to Upper Layer Service from MAC DATA Service

- PDUs are received from the DOCSIS RF.
- If PDU is sent with a suppressed header, the header is regenerated before the packet is subjected to further processing.
- In the CMTS the MAC_DATA service classifies PDUs ingress from the RF using the classification table and then polices the QoS traffic shaping and validates addressing as performed by the CM. In the CM no per-packet service flow classification is required for traffic ingress from the RF.
- Upper layer service receives PDUs from the MAC_DATA.indicate service.

E.4.3 Sample Sequence of MAC Control and MAC Data Services

A possible CM-oriented sequence of MAC service functions for creating, acquiring, modifying, and then using a specific service flow is as follows:

- MAC_REGISTER_RESPONSE.indicate
Learn of any provisioned service flows and their provisioned QoS traffic parameters.
- MAC_CREATE_SERVICE_FLOW.request/response
Create new service flow. This service interface is utilized if the service flow was learned as not provisioned by the MAC_REGISTER_RESPONSE service interface. Creation of a service flow invokes DSA signaling.
- MAC_CHANGE_SERVICE_FLOW.request/response
Define admitted and activated QoS parameter sets, classifiers, and packet suppression headers. Change of a service flow invokes DSC signaling.
- MAC_DATA.request
Send PDUs to MAC service for classification and transmission.

- **MAC_DATA.indication**

Receive PDUs from MAC service.

- **MAC_DELETE_SERVICE_FLOW.request/response**

Delete service flow. Would likely be invoked only for dynamically created service flows, not provisioned service flows. Deletion of a service flow uses DSD signaling.

This page intentionally left blank.

Appendix F. Example Preamble Sequence

F.1 Introduction

A programmable preamble superstring, up to 1024 bits long, is part of the channel-wide profile or attributes, common to the all burst profiles on the channel (Section 6.3.3, Table 6-18), but with each burst profile able to specify the start location within this sequence of bits and the length of the preamble (Section 6.3.3, Table 6-19). The first bit of the Preamble Pattern is designated by the Preamble Value Offset as described in Table 6-19, Section 6.3.3. The first bit of the Preamble Pattern is the first bit into the symbol mapper (Figure 4-8), and is I1 in the first symbol of the burst (see Section 4.2.2.2). As an example, per Table 6-19, for Preamble Offset Value = 100, the 101st bit of the preamble superstring is the first bit into the symbol mapper, and the 102nd bit is the second bit into the mapper, and is mapped to Q1, and so. An example 1024-bit-long preamble superstring is given in Section F.2.

F.2 Example Preamble Sequence

The following is the example 1024-bit preamble sequence:

Bits 1 through 128:

```
1100 1100 1111 0000 1111 1111 1100 0000 1111 0011 1111 0011 0011 0000 0000 1100
0011 0000 0011 1111 1111 1100 1100 1100 1111 0000 1111 0011 1111 0011 1100 1100
```

Bits 129 through 256:

```
0011 0000 1111 1100 0000 1100 1111 1111 0000 1100 1100 0000 1111 0000 0000 1100
0000 0000 1111 1111 1111 0011 0011 0011 1100 0011 1100 1111 1100 1111 0011 0000
```

Bits 257 through 384:

```
1100 0011 1111 0000 0011 0011 1111 1100 0011 0011 0000 0011 1100 0000 0011 0000
0000 1110 1101 0001 0001 1110 1110 0101 0010 0101 0010 0101 1110 1110 0010 1110
```

Bits 385 through 512:

```
0010 1110 1110 0010 0010 1110 1110 1110 1110 0010 0010 0010 1110 1110 0010
1110 1110 1110 0010 1110 0010 1110 0010 0010 0010 0010 1110 0010 0010 1110 0010
```

Bits 513 through 640:

```
0010 0010 1110 1110 1110 1110 1110 0010 1110 0010 1110 0010 1110 1110 0010
0010 1110 1110 0010 1110 1110 1110 0010 1110 1110 0010 1110 0010 0010 1110 0010
```

Bits 641 through 768:

```
0010 1110 1110 1110 0010 0010 0010 1110 0010 1110 1110 1110 0010 0010 1110
0010 1110 0010 0010 0010 1110 1110 0010 0010 0010 0010 1110 0010 0010 0010 0010
```

Bits 769 through 896:

0010 1110 1110 1110 1110 1110 1110 0010 1110 0010 1110 0010 1110 1110 0010 0010
1110 1110 0010 1110 1110 1110 0010 1110 1110 0010 1110 0010 0010 1110 0010 0010

Bits 897 through 1024:

1110 1110 1110 0010 0010 0010 1110 0010 1110 1110 1110 1110 0010 0010 1110 0010
1110 0010 0010 0010 1110 1110 0010 0010 0010 0010 1110 0010 0010 0010 0010 1110

Appendix G DOCSIS v1.0/v1.1 Interoperability

G.1 Introduction

This specification is informally referred to as DOCSIS 1.1. It is the second generation of DOCSIS 1.0 specified in [DOCSIS9]. The terms DOCSIS 1.1 and DOCSIS 1.0 refer to these two different specifications.

The DOCSIS 1.1 specification, primarily aims at enhancing the limited QoS functionality of a DOCSIS 1.0 based cable access system. New MAC messages have been defined for dynamic QoS signaling, and several new QoS parameter encodings have been defined in the existing MAC messages. A DOCSIS 1.1 CMTS can better support the requirements of delay/jitter sensitive traffic on a DOCSIS 1.1 CM.

Besides supporting a rich set of QoS features for DOCSIS 1.1 CMs, the DOCSIS 1.1 CMTS must be backwards compatible with a DOCSIS 1.0 CM. Furthermore, it is necessary for a 1.1 CM to function like a 1.0 CM when interoperating with a 1.0 CMTS.

This section describes the interoperability issues and trade-offs involved, when the operator wishes to support DOCSIS 1.0 as well as DOCSIS 1.1 CMs on the same cable access channel.

G.2 General Interoperability Issues

This section addresses the general DOCSIS 1.0/DOCSIS 1.1 interoperability issues that do not impact the performance during normal operation of the CMs.

G.2.1 Provisioning

The parameters of the TFTP config file for a DOCSIS 1.1 CM, are a superset of those for a DOCSIS 1.0 CM. Configuration file editors will have to be enhanced to incorporate support for these new parameters and the new MIC calculation.

If a DOCSIS 1.1 CM is provisioned with a DOCSIS 1.0 style TFTP configuration file it will register identically to a DOCSIS 1.0 CM. Thus, a DOCSIS 1.1 CM can be provisioned to work seamlessly on either a DOCSIS 1.0 or a DOCSIS 1.1 network. Although, clearly, a DOCSIS 1.1 modem on a DOCSIS 1.0 network would be unable to support any DOCSIS 1.1-specific features.

On the other hand, DOCSIS 1.0 CMs do not recognize (and ignore) many of the new TLVs in a DOCSIS 1.1 style config file, and will be unable to register successfully if provisioned with a DOCSIS 1.1 configuration file. To prevent any functionality mismatches, a DOCSIS 1.1 CMTS MUST reject any Registration Request with DOCSIS 1.1-specific configuration parameters, but without a Modem Capabilities statement that indicates that this is a DOCSIS 1.1 CM.

G.2.2 Registration

A DOCSIS 1.1 CMTS is designed to handle the existing registration TLVs from DOCSIS 1.0 CMs as well as the new TLVs (namely, types 22 to 30) from the DOCSIS 1.1 CM.

There is a slight difference in the Registration related messaging procedure when the DOCSIS 1.1 CMTS is responding to a DOCSIS 1.1 CM as opposed to DOCSIS 1.0 CM. A DOCSIS 1.1 CM could be configured to use the Service Class Name which is statically defined at the CMTS instead of asking for the service class parameters explicitly. When such a Registration-Request is received by the DOCSIS 1.1 CMTS, it encodes the actual parameters of that service class in the Registration-Response and expects the DOCSIS 1.1 specific

Registration-Acknowledge MAC message from the CM. If the detailed capabilities in the Registration-Response message exceed those the CM is capable of supporting, the CM is required to indicate this to the CMTS in its Registration-Acknowledge.

When a DOCSIS 1.0 CM registers with the same CMTS, the default DOCSIS 1.0 version is easily identified by the absence of the "DOCSIS Version" Modem Capabilities encoding in the Registration-Request. The Registration-Request from DOCSIS 1.0 CM explicitly requests all non-default service class parameters in the Registration-Request per its provisioning information. Absence of a Service Class Names eliminates the need for the DOCSIS 1.1 CMTS to explicitly specify the service class parameters in the Registration-Response using DOCSIS 1.1 TLVs. When a DOCSIS 1.1 CMTS receives a Registration-Request from a DOCSIS 1.0 CM, it will respond with the regular DOCSIS 1.0 style Registration-Response and not expect the CM to send the Registration-Acknowledge MAC message.

Another minor issues is that a DOCSIS 1.0 CM will request for a bi-directional (with Upstream/Downstream parameters) service class from the CMTS using a Class-of-Service Configuration Setting.

Since DOCSIS 1.1 CMTS typically operates with unidirectional service classes, it can easily translate a DOCSIS 1.0 Class-of-Service Configuration Setting into DOCSIS 1.1 Service Flow Encodings for setting up unidirectional service classes in local QoS implementation. However, for DOCSIS 1.0 modems, the DOCSIS 1.1 CMTS will continue to maintain the QoSProfile table (with bi-directional Class parameters) for backward compatibility with DOCSIS 1.0 MIB.

Thus, if properly provisioned, a DOCSIS 1.0 and a DOCSIS 1.1 CM can successfully register with the same DOCSIS 1.1 CMTS. Likewise, a DOCSIS 1.0 and a DOCSIS 1.1 CM can successfully register with the same DOCSIS 1.0 CMTS.

G.2.3 Dynamic Service Establishment

There are 8 new MAC messages that relate to Dynamic Service Establishment. A DOCSIS 1.0 CM will never send them to any CMTS since they are unsupported. A DOCSIS 1.1 CM will never send them to a DOCSIS 1.0 CMTS because (a) to register successfully it has to be provisioned as a DOCSIS 1.0 CM and (b) when provisioned as a DOCSIS 1.0 CM it acts identically. When a DOCSIS 1.1 CM is connected to a DOCSIS 1.1 CMTS these messages work as expected.

G.2.4 Fragmentation

Fragmentation is initiated by the CMTS. Thus, a DOCSIS 1.0 CMTS will never initiate fragmentation since it knows nothing about it. A DOCSIS 1.1 CMTS can only initiate fragmentation for DOCSIS 1.1 CMs. A DOCSIS 1.1 CMTS SHOULD NOT attempt to fragment transmissions from DOCSIS 1.0 CMs.

G.2.5 Multicast Support

It is mandatory for DOCSIS 1.0 CM's to support forwarding of multicast traffic. However, the specification is silent on IGMP support. Thus, the only standard mechanism for controlling IP-multicast on DOCSIS 1.0 CMs is through SNMP and packet filters. Designers of DOCSIS 1.0 networks will have to deal with these limitations and expect no different from DOCSIS 1.0 CM's on a DOCSIS 1.1 network.

G.2.6 Upstream Channel Change

A DOCSIS 1.1 CMTS is capable of specifying the level of re-ranging to be performed when it issues an UCC-Request to the CM. This re-ranging technique parameter is specified by the DOCSIS 1.1 CMTS using a new TLV in the UCC-Request MAC message.

DOCSIS 1.1 CMs that recognize this new TLV in the UCC-Request can benefit by only re-ranging to the level specified by this TLV. This can help in reducing the reinitialization time following a UCC, for the DOCSIS 1.1 CM carrying a voice call. A DOCSIS 1.1 CMTS is aware of the type of CM to which it is issuing the UCC-Request. It can refrain from inserting this re-ranging TLV in the UCC-Request for DOCSIS 1.0 CMs. If a DOCSIS 1.1 CMTS inserts this re-ranging TLV in the UCC-Request, the DOCSIS 1.0 CMs which do not recognize this TLV will ignore its contents and perform the default DOCSIS 1.0 re-ranging from start (Initial-Maintenance). The DOCSIS 1.1 CMTS accepts default initial ranging procedure from any modem issued the UCC-Request.

Thus DOCSIS 1.0 and DOCSIS 1.1 CMs on the same upstream channel can be individually requested to change upstream channels without any interoperability issues caused by the DOCSIS 1.1 style re-ranging TLV in the UCC-request.

G.3 Hybrid Devices

Some DOCSIS 1.0 CM designs may be capable of supporting individual DOCSIS 1.1 features via a software upgrade. Similarly, some DOCSIS 1.0 CMTS's may be capable of supporting individual DOCSIS 1.1 features.¹ To facilitate these "hybrid" devices, the majority of DOCSIS 1.1 features are individually enumerated in the Modem Capabilities.

DOCSIS 1.0 hybrid CM's MAY request DOCSIS 1.1 features via this mechanism. However, unless a CM is fully DOCSIS 1.1 compliant (i.e. not a hybrid), it MUST NOT send a "DOCSIS Version" Modem Capability which indicates anything besides DOCSIS 1.0.

If a hybrid CM intends to request such 1.1 capabilities from the CMTS during registration, it MUST send the ASCII coded string in Option code 60 of its DHCP request, "docsis1.0:xxxxxxx". Where xxxxxxxx must be the hexadecimal encoding of the Modem Capabilities, refer to C.1.3.1. The DHCP server MAY use such information to determine what configuration file the CM is to use.²

Normally, a DOCSIS 1.0 CMTS would set all unknown Modem Capabilities to 'Off' in the Registration Response indicating that these features are unsupported and MUST NOT be used by the CM. A DOCSIS 1.0 hybrid CMTS's MAY leave supported Modem Capabilities set to 'On' in the Registration Response. However, unless a CMTS is fully DOCSIS 1.1 compliant (i.e. not a hybrid), it MUST still set all "DOCSIS Version" Modem Capabilities to DOCSIS 1.0.

As always, any Modem Capability set to 'Off' in the Registration Response must be viewed as unsupported by the CMTS and MUST NOT be used by the CM.

G.4 Interoperability & Performance

This section addresses the issue of performance impact on the QoS for DOCSIS 1.1 CMs when DOCSIS 1.0 and DOCSIS 1.1 CMs are provisioned to share the same upstream MAC channel.

The DOCSIS 1.0 CMs lack the ability to explicitly set their request policy (or provide scheduling parameters) for the advanced DOCSIS 1.1 scheduling mechanisms like "Unsolicited Grant Service" and "Real-Time Polling Service". Thus, DOCSIS 1.0 CMs will only receive statically configured "Tiered Best Effort" or "CIR" service on the upstream. The DOCSIS 1.1 CMs on the same upstream channel can explicitly request for additional

1. Section G.3, par.1, edited 06/21/99 per rfi-n-99051

2. Section G.3, paragraph 3 added 6/9/99 per ECN rfi-n-99025

Service Flows when required, using the DOCSIS 1.1 DSA-Request MAC message. Thus, DOCSIS 1.1 CMs can benefit from the advanced scheduling mechanisms of a DOCSIS 1.1 CMTS for their real-time traffic, besides the best-effort scheduling service they share with the DOCSIS 1.0 CMs on the same upstream channel.

The DOCSIS 1.1 upstream cable access channel carries variable-length MAC frames. In spite of the variable-length nature of the MAC frames, the DOCSIS 1.1 CMTS grant scheduler is theoretically capable of providing a zero jitter TDMA-like environment for voice grants on the Upstream. Whenever the grant scheduler detects that the deadline of any future voice grant will be violated by the insertion of a non-voice grant, it fragments the non-voice grant up to the future voice grant boundary. Thus the voice grants see a zero shift from the assigned periodic grant position.

However, such grant fragmentation might not always be possible when the CMTS supports DOCSIS 1.0 CMs along with DOCSIS 1.1 CMs on the same Upstream channel since DOCSIS 1.0 CM do not support fragmentation. For a mixed CM version upstream channel, the worst case voice grant jitter seen by the DOCSIS 1.1 CMs, is when a DOCSIS 1.0 CM is given a grant for an unfragmented maximum sized MAC frame just before the designated voice grant slot of the DOCSIS 1.1 CM.

The maximum Voice grant jitter experienced by the DOCSIS 1.1 CMs is a function of the physical layer characteristics of the Upstream Channel. For 10.24Mbps and 5.12Mbps upstream channels, the impact of having fragmenting and non-fragmenting CM's on the same channel is almost undetectable. On smaller channels, the benefit of fragmentation is far greater and the jitter induced by non-fragmenting DOCSIS 1.0 CM's is greater.

Thus, properly engineered networks can support voice even when mixing DOCSIS 1.0 and DOCSIS 1.1 CM's.

Appendix H. Multiple Upstream Channels

This appendix presents an example of several upstream channels served by a single downstream channel. This is meant to illustrate one topology and one implementation of that topology.

Suppose one downstream channel is used in conjunction with four upstream channels as shown in Figure H-1. In this case, the four upstream channels are separate fibers serving four geographical communities of modems.

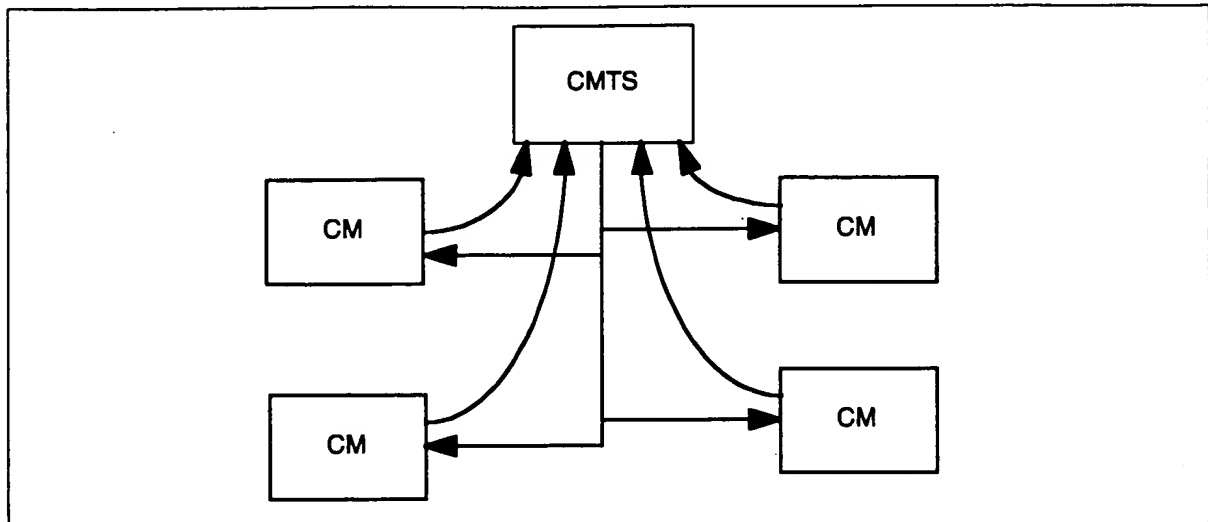


Figure H-1. One Downstream and Four Upstream Channels

In this topology, the CMTS transmits four Upstream Channel Descriptors (UCDs) and four MAPs. Unfortunately, each CM cannot determine to which upstream channel it is attached, because there is no way to convey the geographical information on the shared downstream channel. The CM must assume (at least at initialization) that the UCD and MAP apply to the channel to which it is attached. The CM chooses an Initial Maintenance opportunity on any of the channels and transmits a Ranging Request. The CMTS will receive the request and will redirect the CM to the appropriate upstream channel identifier. From then on, the CM will be using the MAP that is appropriate to the fiber branch to which it is connected.

A number of constraints are imposed by this topology:

- All of the upstream channels must operate at the same frequency. Since the CM is choosing a channel descriptor at random, it would be transmitting on the wrong frequency if it chose the UCD that applied to a different fiber path.
- All of the upstream channels must operate at the same symbol rate. If not, the CMTS would be unable to demodulate the Ranging Request if transmitted at the wrong symbol rate for the particular channel.
- All Initial Maintenance opportunities across all fiber branches must be aligned. When the CM randomly chooses a MAP to use, the CMTS must be prepared to receive a Ranging Request at that time.
- All Initial Maintenance opportunities must use the same burst characteristics so that the CMTS can demodulate the Ranging Request.

Note that only the initialization intervals must be aligned. Once the CM is assigned its proper channel ID, its activities need only be aligned with other users of its fiber branch. Ordinary data transmission and requests for bandwidth may occur independently across the four upstream channels.

This page intentionally left blank.

Appendix I. The Data-Over-Cable Spanning Tree Protocol

Section 3.1.2.1 requires the use of the spanning tree protocol on CMs that are intended for commercial use and on bridging CMTSs. This appendix describes how the 802.1d spanning tree protocol is adapted to work for data over cable systems.

I.1 Background

A spanning tree protocol is frequently employed in a bridged network in order to deactivate redundant network connections; i.e., to reduce an arbitrary network mesh topology to an active topology that is a rooted tree that spans all of the network segments. The spanning tree algorithm and protocol should not be confused with the data-forwarding function itself; data forwarding may follow transparent learning bridge rules, or may employ any of several other mechanisms. By deactivating redundant connections, the spanning tree protocol eliminates topological loops, which would otherwise cause data packets to be forwarded forever for many kinds of forwarding devices.

A standard spanning tree protocol [IEEE 802.1d] is employed in most bridged local area networks. This protocol was intended for private LAN use and requires some modification for cable data use.

I.2 Public Spanning Tree

To use a spanning tree protocol in a public-access network such as data-over-cable, several modifications are needed to the basic IEEE 802.1d process. Primarily, the public spanning tree must be isolated from any private spanning tree networks to which it is connected. This is to protect both the public cable network and any attached private networks. Figure I-1 illustrates the general topology.

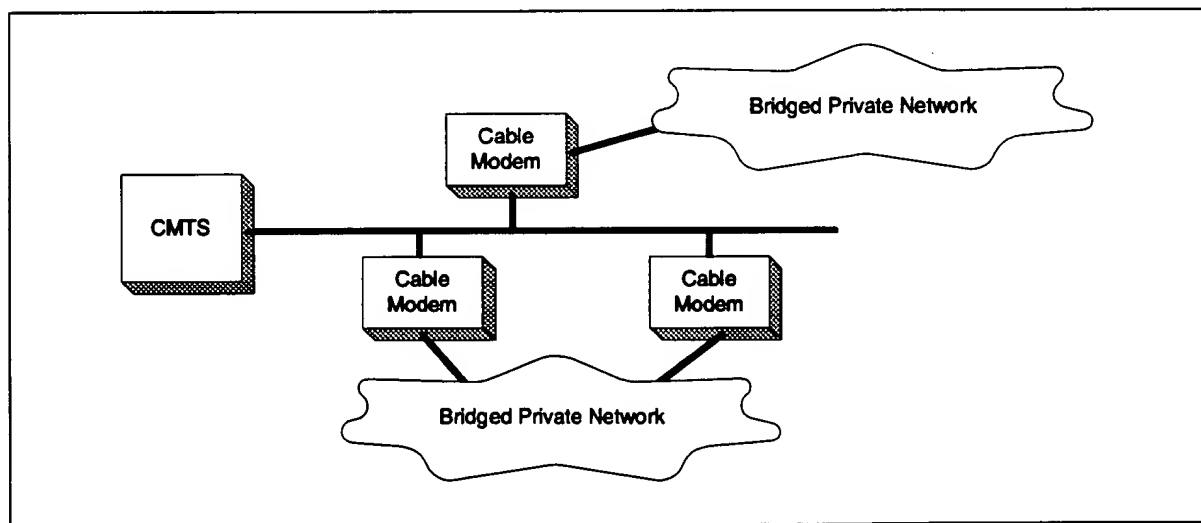


Figure I-1. Spanning Tree Topology

The task for the public spanning tree protocol, with reference to Figure I-1, is to:

- Isolate the private bridged networks from each other. If the two private networks merge spanning trees then each is subject to instabilities in the other's network. Also, the combined tree may exceed the maximum allowable bridging diameter.

- Isolate the public network from the private networks' spanning trees. The public network must not be subject to instabilities induced by customers' networks; nor should it change the spanning tree characteristics of the customers' networks.
- Disable one of the two redundant links into the cable network, so as to prevent forwarding loops. This should occur at the cable modem, rather than at an arbitrary bridge within the customer's network.

The spanning tree protocol must also serve the topology illustrated in Figure I-2:




Figure I-2. Spanning Tree Across CMTSs

In Figure I-2, in normal operation the spanning tree protocol should deactivate a link at one of the two cable modems. It should not divert traffic across the private network. Note that in some circumstances, such as deactivation of Link-X, spanning tree *will* divert traffic onto the private network (although limits on learned MAC addresses will probably throttle most transit traffic). If this diversion is undesirable, then it must be prevented by means external to spanning tree; for example, by using routers.

I.3 Public Spanning Tree Protocol Details

The Data over Cable Spanning Tree algorithm and protocol is identical to that defined in [IEEE 802.1d], with the following exceptions:

- When transmitting Configuration Bridge Protocol Data Units (BPDUs), the Data over Cable Spanning Tree Multicast Address 01-E0-2F-00-00-03 is used rather than that defined in IEEE 802.1d. These BPDUs will be forwarded rather than recalculated by ordinary IEEE 802.1d bridges.
- When transmitting Configuration BPDUs, the SNAP header AA-AA-03-00-E0-2F-73-74 is used rather than the LLC 42-42-03 header employed by 802.1d. This is to further differentiate these BPDUs from those used by IEEE 802.1d bridges, in the event that some of those bridges do not correctly identify multicast MAC addresses.¹
- IEEE 802.1d BPDUs are ignored and silently discarded.

1. It is likely that there are a number of spanning tree bridges deployed which rely solely on the LSAPs to distinguish 802.1d packets. Such devices would not operate correctly if the data-over-cable BPDUs also used LSAP=0x42.

- Topology Change Notification (TCN) PDUs are not transmitted (or processed). TCNs are used in IEEE networks to accelerate the aging of the learning database when the network topology may have changed. Since the learning mechanism within the cable network typically differs, this message is unnecessary and may result in unnecessary flooding.
- CMTSs operating as bridges must participate in this protocol and must be assigned higher priorities (more likely to be root) than cable modems. The NSI interface on the CMTS SHOULD be assigned a port cost equivalent to a link speed of at least 100 Mbps. These two conditions, taken together, should ensure that (1) a CMTS is the root, and (2) any other CMTS will use the head-end network rather than a customer network to reach the root.
- The MAC Forwarder of the CMTS MUST forward BPDUs from upstream to downstream channels, whether or not the CMTS is serving as a router or a bridge.

Note that CMs with this protocol enabled will transmit BPDUs onto subscriber networks in order to identify other CMs on the same subscriber network. These public spanning tree BPDUs will be carried transparently over any bridged private subscriber network. Similarly, bridging CMTSs will transmit BPDUs on the NSI as well as on the RFI interface. The multicast address and SNAP header defined above are used on all links.

1.4 Spanning Tree Parameters and Defaults

Section 4.10.2 of [IEEE 802.1d] specifies a number of recommended parameter values. Those values should be used, with the exceptions listed below:

Path Cost

In [IEEE 802.1d], the following formula is used:

$$\text{Path_Cost} = 1000 / \text{Attached_LAN_speed_in_Mb/s}$$

For CMs, this formula is adapted as:

$$\text{Path_Cost} = 1000 / (\text{Upstream_symbol_rate} * \text{bits_per_symbol_for_long_data_grant})$$

That is, the modulation type (QPSK or 16QAM) for the Long Data Grant IUC is multiplied by the raw symbol rate to determine the nominal path cost. Table I-1 provides the derived values.

Table I-1. CM Path Cost

Symbol Rate	Default Path Cost	
	QPSK	16QAM
ksym/sec		
160	3125	1563
320	1563	781
640	781	391
1280	391	195
2560	195	98

For CMTSs, this formula is:

$$\text{Path_Cost} = 1000 / (\text{Downstream_symbol_rate} * \text{bits_per_symbol})$$

Bridge Priority

The Bridge Priority for CMs SHOULD default to 36864 (0x9000). This is to bias the network so that the root will tend to be at the CMTS. The CMTS SHOULD default to 32768, as per 802.1d.

Note that both of these recommendations affect only the *default* settings. These parameters, as well as others defined in 802.1d, SHOULD be manageable throughout their entire range through the Bridge MIB (RFC-1493) or other means.

Appendix J. Error Codes and Messages¹

These are CM and CMTS error codes and messages. These error codes are meant to emulate the standard fashion that ISDN reports error conditions regardless of the vendor producing the equipment.

The errors reported are Sync loss, UCD, MAP, Ranging REQ/RSP, UCC, registration, dynamic service request, and DHCP/TFTP failures. In some cases there is detailed error reports in other error codes are simply "it failed."

Table J-1. Error Codes for MAC Management Messages

Error Code	Error Message
T00.0	SYNC Timing Synchronization
T01.0	Failed to acquire QAM/QPSK symbol timing. Error stats? Retry #'s?
T02.0	Failed to acquire FEC framing. Error stats? Retry #'s? # of bad frames?
T02.1	Acquired FEC framing. Failed to acquire MPEG2 Sync. Retry #'s?
T03.0	Failed to acquire MAC framing. Error stats? Retry #'s? # of bad frames?
T04.0	Failed to Receive MAC SYNC frame within time-out period.
T05.0	Loss of Sync. (Missed 5 in a row, after having SYNC'd at one time)
U00.0	UCD Upstream Channel Descriptor
U01.0	No UCD's Received. Time-out.
U02.0	UCD invalid or channel unusable.
U03.0	UCD valid, BUT no SYNC received. TIMED OUT.
U04.0	UCD, & SYNC valid, NO MAPS for THIS Channel.
U05.0	UCD received with invalid or out of order Configuration Change Count.
U06.0	US Channel wide parameters not set before Burst Descriptors.
M00.0	MAP Upstream Bandwidth Allocation
M01.0	A transmit opportunity was missed because the MAP arrived too late.
R00.0	RNG-REQ Ranging Request
R01.0	NO Maintenance Broadcasts for Ranging opportunities Received T2 time-out.
R04.0	Received Response to Broadcast Maintenance Request, But no Unicast Maintenance opportunities received. T4 time-out.
R101.0	No Ranging Requests received from POLLED CM (CMTS generated polls).
R102.0	Retries exhausted for polled CM (report MAC address). After 16 R101.0 errors.
R103.0	Unable to Successfully Range CM (report MAC address) Retries Exhausted. Note: this is different from R102.0 in that it was able to try, i.e. got REQ's but failed to Range properly.
R104.0	Failed to receive Periodic RNG-REQ from modem (SID X), timing-out SID.
R00.0	RNG-RSP Ranging Response

1. Appendix entirely replaced per rfi-n-99049 06/30/99.ew

R02.0	No Ranging Response received, T3 time-out.
R03.0	Ranging Request Retries exhausted.
R05.0	Started Unicast Maintenance Ranging no Response received. T3 time-out.
R06.0	Unicast Maintenance Ranging attempted. No Response. Retries exhausted.
R07.0	Unicast Ranging Received Abort Response. Re-initializing MAC.
I00.0	REG-REQ Registration Request
I04.0	Service not available. Reason: Other.
I04.1	Service not available. Reason: Unrecognized configuration setting.
I04.2	Service not available. Reason: Temporarily unavailable.
I04.3	Service not available. Reason: Permanent.
I101.0	Invalid MAC header.
I102.0	Invalid SID, not in use.
I103.0	Required TLV's out of order.
I104.0	Required TLV's not present.
I105.0	Down Stream Frequency format invalid.
I105.1	Down Stream Frequency not in use.
I105.2	Down Stream Frequency invalid, not a multiple of 62500Hz.
I106.0	Up Stream Channel invalid, unassigned.
I106.1	Up Stream Channel Change followed with (RE-)Registration REQ.
I107.0	Up Stream Channel overloaded.
I108.0	Network Access configuration has invalid parameter.
I109.0	Class of Service configuration is invalid.
I110.0	Class of Service ID unsupported.
I111.0	Class of Service ID invalid or out of range.
I112.0	Max Down Stream Bit Rate configuration is invalid format.
I112.1	Max Down Stream Bit Rate configuration setting is unsupported.
I113.0	Max Up Stream Bit Rate configuration setting invalid format.
I113.1	Max Up Stream Bit Rate configuration setting unsupported.
I114.0	Up Stream Priority configuration invalid format.
I114.1	Up Stream Priority configuration setting out of range.
I115.0	Guaranteed Min Up Stream Channel Bit Rate configuration setting invalid format.
I115.1	Guaranteed Min Up Stream Channel Bit Rate configuration setting exceeds Max Up Stream Bit Rate.
I115.2	Guaranteed Min Up Stream Channel Bit Rate configuration setting out of range.
I116.0	Max Up Stream Channel Transmit Burst configuration setting invalid format.
I116.1	Max Up Stream Channel Transmit Burst configuration setting out of range.
I117.0	Modem Capabilities configuration setting invalid format.
I117.1	Modem Capabilities configuration setting.
I200.0	Version 1.1 Specific REG-REQ Registration Request
I201.0	Registration rejected unspecified reason.
I201.1	Registration rejected unrecognized configuration setting.
I201.2	Registration rejected temporary no resource.
I201.3	Registration rejected permanent administrative.
I201.4	Registration rejected required parameter not present.
I201.5	Registration Rejected header suppression setting not supported.

100.0	REG-RSP Registration Response
I01.0	Registration RESP invalid format or not recognized.
I02.0	Registration RESP not received.
I03.0	Registration RESP with bad SID.
I300.0	REG-ACK Registration Acknowledgement
I301.0	Registration aborted no REG-ACK.
C00.0	UCC-REQ Upstream Channel Change Request
C01.0	UCC-REQ received with invalid or out of range US channel ID.
C02.0	UCC-REQ received unable to send UCC-RSP, no TX opportunity.
C100.0	UCC-RSP Upstream Channel Change Response
C101.0	UCC-RSP not received on previous channel ID.
C102.0	UCC-RSP received with invalid channel ID.
C103.0	UCC-RSP received with invalid channel ID on new channel.
D00.0	DHCP CM Net Configuration download and Time of Day
D01.0	Discover sent no Offer received, No available DHCP Server.
D02.0	Request sent, no Response.
D03.0	Requested Info not supported.
D03.1	DHCP response doesn't contain ALL the valid fields as describe in the RF spec Appendix D
D04.0	Time of Day, none set or invalid data.
D04.1	Time of Day Request sent no Response received.
D04.2	Time of Day Response received but invalid data/format.
D05.0	TFTP Request sent, No Response/No Server.
D06.0	TFTP Request Failed, configuration file NOT FOUND.
D07.0	TFTP Failed, OUT OF ORDER packets.
D08.0	TFTP complete, but failed Integrity Check (MIC).
S00.0	Dynamic Service Requests
S01.0	Service add rejected unspecified reason.
S01.1	Service add rejected unrecognized configuration setting.
S01.2	Service add rejected temporary no resource.
S01.3	Service add rejected permanent administrative.
S01.4	Service add rejected required parameter not present.
S01.5	Service add rejected header suppression setting not supported.
S01.6	Service add rejected service flow exists.
S01.7	Service add rejected HMAC authentication failure.
S02.0	Service change rejected unspecified reason.
S02.1	Service change rejected unrecognized configuration setting.
S02.2	Service change rejected temporary no resource.
S02.3	Service change rejected permanent administrative.
S02.4	Service change rejected requestor not owner of service flow.
S02.5	Service change rejected service flow not found.

S02.6	Service change rejected required parameter not present.
S02.5	Service change rejected header suppression setting not supported.
S02.6	Service change rejected HMAC authentication failure.
S03.0	Service delete rejected unspecified reason.
S03.1	Service delete rejected requestor not owner of service flow.
S03.2	Service delete rejected service flow not found.
S03.3	Service delete rejected HMAC authentication failure.
S100.0	Dynamic Service Responses
S101.0	Service add response rejected invalid transaction ID.
S102.0	Service change response rejected invalid transaction ID.
S103.0	Service delete response rejected invalid transaction ID.
S200.0	Dynamic Service Acknowledgements
S201.0	Service add ACK rejected invalid transaction ID.
S201.1	Service add aborted no ACK.
S202.0	Service change ACK rejected invalid transaction ID.
S202.1	Service change aborted no ACK.
B00.0	Baseline Privacy
B01.0	TBD

Appendix K. DOCSIS Transmission and Contention Resolution

K.1 Introduction:

This Appendix attempts to clarify how the DOCSIS transmission and contention resolution algorithms work. It has a few minor simplifications and a few assumptions, but should definitely help clarify this area of the spec.

This example has a few simplifications:

- It doesn't explicitly talk about packet arrivals while deferring or waiting for pending grants and is vague about sizing piggyback requests.
- Much of this applies with concatenation, but it does not attempt to address all the subtleties of that situation.

It also has a few assumptions:

- It assumes that a Request always fits in any Request/Data region.
- When a piggyback request is sent with a contention data packet, the state machine only checks for the Grant to the Request and assumes the Data Ack for the contention data packet was supplied by CMTS.
- It probably assumes a few other things, but should be sufficient to get the basic point across.

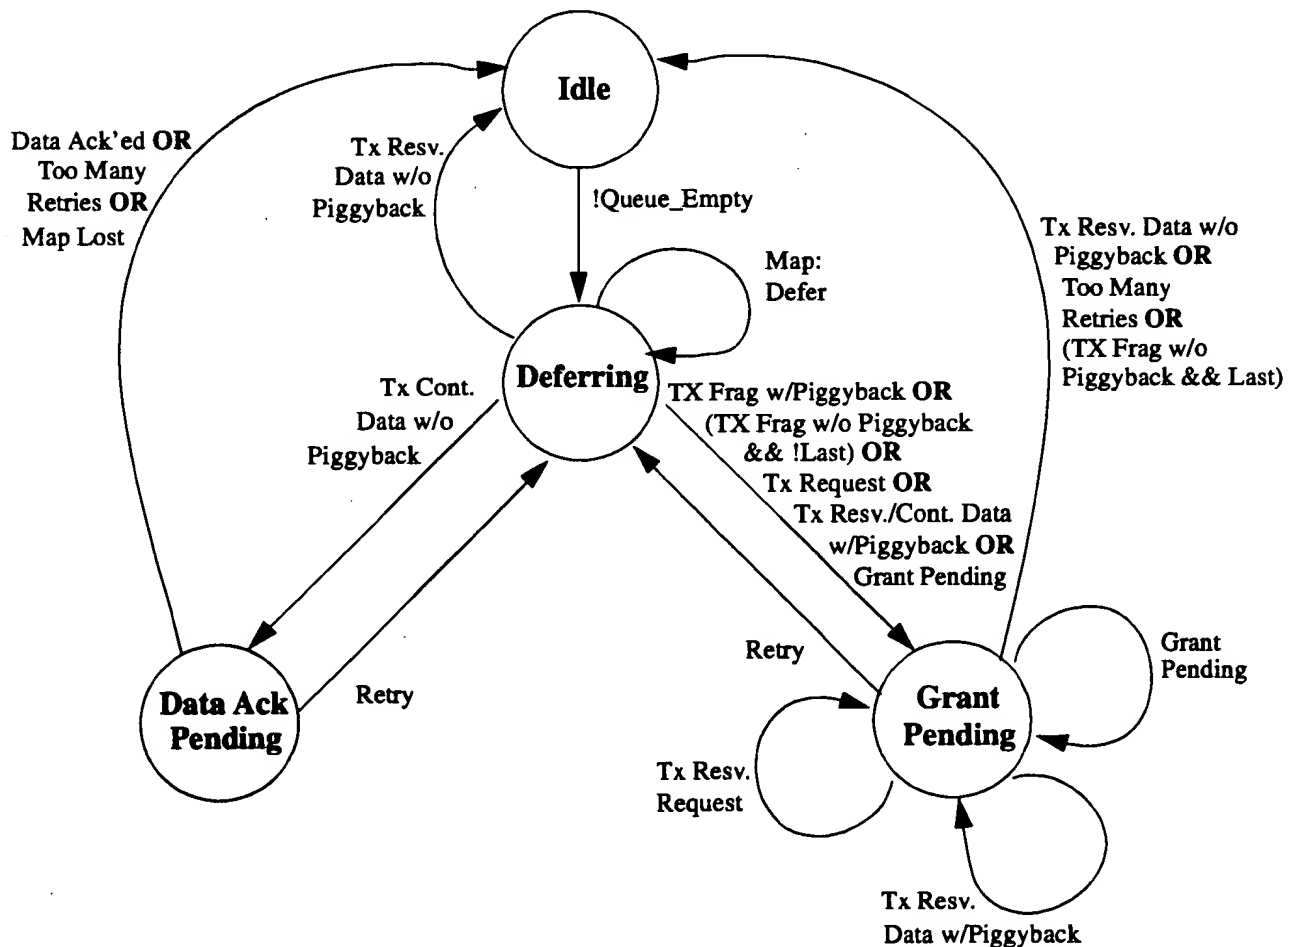


Figure K-1. Transmission & Deference State Transition Diagram

Start	= Data Backoff Start field from Map “currently in effect”
End	= Data Backoff End field from Map “currently in effect”
Window	= Current backoff window
Random[n]	= Random number generator that selects a number between 0 and n-1
Defer	= Number of Transmit Opportunities to defer before transmitting
Retries	= Number of transmissions attempted without resolution
Tx_time	= Saved time of when Request or Request/Data was transmitted
Ack_time	= Ack Time field from current Map
Piggyback	= Flag set whenever a piggyback REQ is added to a transmit pkt
Queue_Empty	= Flag set whenever the data queue for this SID is empty
Lost_Map	= Flag set whenever a MAP is lost & we’re in state Data Ack Pending
my_SID	= Service ID of the queue that has a packet to transmit
pkt size	= Data packet size including MAC and physical layer overhead (including piggyback if used)
frag_size	= Size of the fragment
Tx_Mode	= {Full_Pkt; First_Frag; Middle_Frag; Last_Frag}
min_frag	= Size of the minimum fragment

```
Window = 0;
Retries = 0;
```

CalcDefer();
go to Deferring

Wait for next Map;

```

if (Data_Acknowledge_SID == my_SID)      /* Success! CMTS received data packet */
    go to state Idle;
else if (Ack_time > Tx_time)              /* COLLISION!!! or Pkt Lost or Map Lost */
{
    if (Lost_Map)
        go to state Idle;                /* Assume pkt was ack'ed to avoid sending duplicates */
    else
        Retry();
}

```

stay in state Data Ack Pending;

Wait for next Map;

```
while (Grant SID == my_SID)
    UtilizeGrant();
```

```

if (Ack_time > Tx_time)          /* COLLISION!!!! or Request denied/lost or Map Lost */
    Retry();
stay in state Grant Pending

```

State: Deferring — Determine Proper Transmission Timing & Transmit

```

if (Grant SID == my_SID)                                /* Unsolicited Grant */
{
    UtilizeGrant();
}
else if (unicast Request SID == my_SID)                  /* Unsolicited Unicast Request */
{
    transmit Request in reservation;
    Tx_time = time;

    go to state Grant Pending;
}
else
{
    for (each Request or Request/Data Transmit Opportunity)
    {
        if (Defer != 0)
            Defer = Defer - 1;                            /* Keep deferring until Defer = 0 */
        else
        {
            if (Request/Data tx_op) and (Request/Data size >= pkt size) /* Send data in contention */
            {
                transmit data pkt in contention;
                Tx_time = time;

                if (Piggyback)
                    go to state Grant Pending;
                else
                    go to state Data Ack Pending;
            }
            else                                           /* Send Request in contention */
            {
                transmit Request in contention;
                Tx_time = time;

                go to state Grant Pending;
            }
        }
    }
}

```

Wait for next Map;
stay in state Deferring

Function: CalcDefer() — Determine Defer Amount

```

if (Window < Start)
    Window = Start;

if (Window > End)
    Window = End;

Defer = Random[2^Window];

```

Function: UtilizeGrant() — Determine Best Use of a Grant

```

if (Grant size >= pkt size)                /* CM can send full pkt */
{
    transmit packet in reservation;
    Tx_time = time;
    Tx_mode = Full_pkt

    if (Piggyback)
        go to state Grant Pending
    else
        go to state Idle;
}
else if (Grant size < min_frag && Grant Size > Request size) /* Can't send fragment, but can send a Request */
{
    transmit Request in reservation;
    Tx_time = time;

    go to state Grant Pending;
}
else if (Grant size == 0)                   /* Grant Pending */
    go to state Grant Pending;
else
{
    while (pkt_size > 0 && Grant SID == my_SID)
    {
        if (Tx_mode == Full_Pkt)
            Tx_mode = First_frag;
        else
            Tx_mode = Middle_frag;

        pkt_size = pkt_size - frag_size;
        if (pkt_size == 0)
            Tx_mode = Last_frag;

        if (another Grant SID == my_SID)    /* multiple grant mode */
            piggyback_size = 0
        else
            piggyback_size = pkt_size       /* piggyback mode */

        if (piggyback_size > 0)
            transmit fragment with piggyback request for remainder of packet in reservation
        else
            transmit fragment in reservation;
    }

    go to state Grant Pending;
}

```

Function: Retry()

Retries = Retries + 1;

if (Retries > 16)

```
{
    discard pkt, indicate exception condition
    go to state Idle;
}
```

Window = Window + 1;

CalcDefer();

go to state Deferring;

This page intentionally left blank.

Appendix L IGMP Example

Section 3.3.1 defines the requirements for CMTS and CM support of IGMP signalling. This Appendix provides further details on CM support for IGMP.

The process defined MAY be supported by compliant CMs. Refer to Figure L-1.

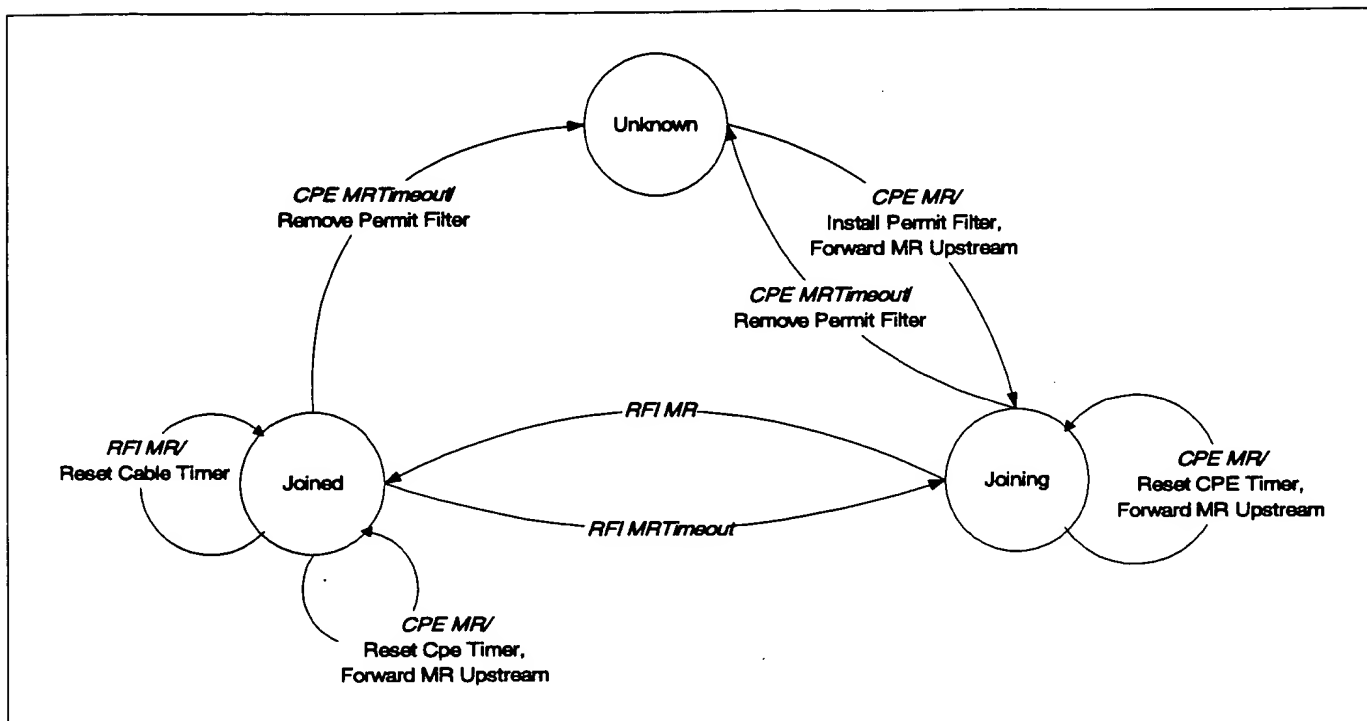


Figure L-1. IGMP Support - CM

Table L-1. Event Table

Event	State		
	1. Unknown	2. Joining	3. Joined
A. CpeMR	Joining	Joining	Joined
B. RFI MR		Joined	Joined
C. RFI MRTIMEOUT			Joining
E. CpeMRTIMEOUT		Unknown	Unknown

L.1 Transition Events

1A

- Forward Membership Report (MR) Upstream
- Start CPE MR Timer
- Install Permit Multicast Filters for forwarding IP multicast traffic to the CPE LAN

2A

- Restart CPE MR timer
- Forward MR upstream

3A

- Stop Cable MR timer

1B

- Start Cable MR timer

2B

- Start Cable MR timer

3B

- Restart Cable MR timer

3C

- Stop Cable MR timer

2D

- Stop CPE MR timer
- Remove Permit Multicast Filter for forwarding IP multicast to the CPE LAN

3D

- Stop CPE MR timer
- Remove Permit Multicast Filter for forwarding IP multicast to the CPE LAN

Appendix M. Unsolicited Grant Services¹

This appendix discusses the intended use of the Unsolicited Grant Service (UGS) and Unsolicited Grant Service with Activity Detection (UGS-AD) and includes specific examples.

M.1 Unsolicited Grant Service (UGS)

M.1.1 Introduction

Unsolicited Grant Service is an Upstream Flow Scheduling Service Type that is used for mapping constant bit rate (CBR) traffic onto Service Flows. Since the upstream is scheduled bandwidth, a CBR service can be established by the CMTS scheduling a steady stream of grants. These are referred to as unsolicited because the bandwidth is predetermined, and there are no ongoing requests being made.

The classic example of a CBR application of interest is Voice over Internet Protocol (VoIP) packets. Other applications are likely to exist as well.

Upstream Flow Scheduling Services are associated with Service Flows, each of which is associated with a single Service ID (SID). Each Service Flow may have multiple Classifiers. Each Classifier may be associated with a unique CBR media stream. Classifiers may be added and removed from a Service Flow. Thus, the semantics of UGS must accommodate single or multiple CBR media streams per SID.

For the discussion within this Appendix, a Subflow will be defined as the output of a Classifier. Since a VoIP session is identified with a Classifier, a Subflow in this context refers to a VoIP session.

M.1.2 Configuration Parameters

- Nominal Grant Interval
- Unsolicited Grant Size
- Tolerated Grant Jitter
- Grants per Interval

Explanation of these parameters and their default values are provided in Appendix C.

M.1.3 Operation

When a Service Flow is provisioned for UGS, the Nominal Grant Interval is chosen to equal the packet interval of the CBR application. For example, VoIP applications with 10 ms packet sizes will require a Nominal Grant Interval of 10 ms. The size of the grant is chosen to satisfy the bandwidth requirements of the CBR application and relates directly to the length of the packet.

When multiple Subflows are assigned to a UGS service, multiple grants per interval are issued. There is no explicit mapping of Subflows to grants. The multiple grants per interval form a pool of grants in which any subflow can use any grant.

It is assumed in this operational example the default UGS case of no concatenation and no fragmentation.

1. Section renamed per rfi-n-99043, item b. ew

M.1.4 Jitter

Figure M-1 shows the relationship between Grant Interval and Tolerated Grant Jitter, and shows an example of jitter on subflows.

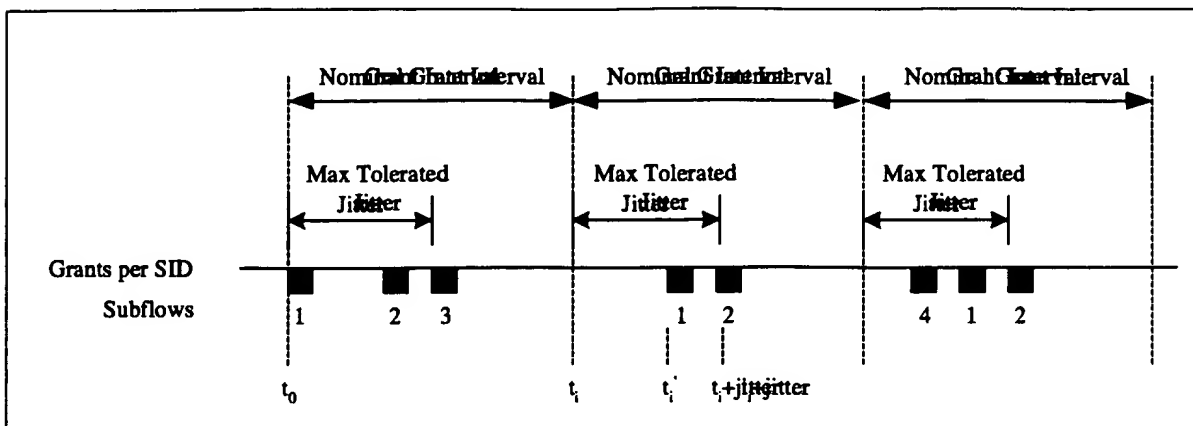


Figure M-1. Example Jitter with Multiple Grants per SID

For only one Grant per Interval, the Tolerated Grant Jitter is the maximum difference between the actual grant time (t_1') and the nominal grant time (t_1). For multiple Grants per Interval, the Tolerated Grant Jitter is the maximum difference between the actual time of the last grant in the group of grants and the nominal grant time (t_1). If the arrival of any grant is at t_1' , then $t_1 \leq t_1' \leq t_1 + \text{jitter}$.

Figure M-1 demonstrates how a Subflow will be jittered even though the individual grants may not move from their relative position. During the first interval, three VoIP sessions are established, and they happen fall on the three grants. In the second interval, VoIP session 3 has been torn down. Since the CMTS does not know which Subflow is associated with which grant, it decides to remove the first grant. The remaining two calls shift to the other two grants. In the third interval, a new VoIP session 4 and a new grant have been added. The new call happens to fall on the new grant. The net effect is that the Subflows may move around within their jitter interval.

The advantage of a small jitter interval is that the VoIP receive jitter buffer may be kept small. The disadvantage is that this places a scheduling constraint on the CMTS.

The boundary of a Nominal Grant Interval is arbitrary and is not communicated between the CMTS and the CM.

Note: More dramatic events like the loss of a downstream MAP, or the frequency hopping of an upstream may cause subflows to jitter outside of this jitter window.

M.1.5 Synchronization Issues¹

There are two synchronization problems that occur when carrying CBR traffic such as VoIP sessions across a network. The first is a frequency mismatch between the source clock and the destination clock. This is managed by the VoIP application, and is beyond the scope of this specification. The second is the frequency mismatch between the CBR source/sinks, and the bearer channel that carries them.

Specifically, if the clock that generates the VoIP packets towards the upstream is not synchronized with the clock at the CMTS which is providing the UGS service, the VoIP packets may begin to accumulate in the CM. This could also occur if a MAP was lost, causing packets to accumulate.

1. Title changed per rfi-r-99043 06/21/99 ew

When the CM detects this condition, it asserts the Queue Indicator in the Service Flow EH Element. The CMTS will respond by issuing an occasional extra grant so as to not exceed 1% of the provisioned bandwidth. (This corresponds to a maximum of one extra grant every one hundred grants). The CMTS will continue to supply this extra bandwidth until the CM deasserts this bit.

A similar problem occurs in the downstream. The far end transmitting source may not be frequency synchronized to the clock which drives the CMTS. Thus the CMTS should police at a rate slightly higher than the exact provisioned rate to allow for this mismatch and to prevent delay buildup or packet drops at the CMTS.

M.2 Unsolicited Grant Service with Activity Detection (UGS-AD)

M.2.1 Introduction

Unsolicited Grant Service with Activity Detection (UGS-AD) is an Upstream Flow Scheduling Service Type. This section describes one application of UGS-AD which is the support for Voice Activity Detection (VAD). VAD is also known as Silence Suppression and is a voice technique in which the transmitting CODEC sends voice samples only when there is significant voice energy present. The receiving CODEC will compensate for the silence intervals by inserting silence or comfort noise equal to the perceived background noise of the conversation.

The advantage of VAD is the reduction of network bandwidth required for a conversation. It is estimated that 60% of a voice conversation is silence. With that silence removed, that would allow a network to handle substantially more traffic.

Subflows in this context will be described as active and inactive. Both of these states of within the MAC Layer QOS state known as Active.

M.2.2 MAC Configuration Parameters

The configuration parameters include all of the normal UGS parameters, plus:

- Nominal Polling Interval
- Tolerated Poll Jitter

Explanation of these parameters and their default values are provided in <Appendix C>.

M.2.3 Operation

When there is no activity, the CMTS sends polled requests to the CM. When there is activity, the CMTS sends Unsolicited Grants to the CM. The CM indicates in the UGS_Parm field of the Service Flow EH Element in each packet of each Unsolicited Grant the number of Grants per Interval that it currently requires. The CM may request up to the maximum active Grants per Interval. The CM constantly sends this state information so that no explicit acknowledgment is required from the CMTS.

It is left to the implementation of the CM to determine activity levels. Implementation options include:

- Having the MAC layer service provide an activity timer per Classifier. The MAC layer service would mark a Subflow inactive if packets stopped arriving for a certain time, and mark a Subflow active the moment a new packet arrived. The number of Grants requested would equal the number of active Subflows.
- Having a higher layer service entity such as an embedded media client which indicates activity to the MAC layer service.

When the CM is receiving polled requests and it detects activity, the CM requests enough bandwidth for one Grant per Interval. If activity is for more than one Subflow, the CM will indicate this in the UGS_Parm field of the Service Flow EH Element beginning with the first packet it sends.

When the CM is receiving Unsolicited Grants, then detects new activity, and asks for one more grant, there will be a delay in time before it receives the new grant. During that delay, packets may build up at the CM. When the new Unsolicited Grant is added, the CMTS will burst extra Grants to clear out the packet buildup.

When the CM is receiving Unsolicited Grants, then detects inactivity on a Subflow and asks for one less grant, there will be a delay in time before the reduction in Grants occurs. If there has been any build up of packets in the upstream transmit queue, the extra grants will reduce or empty the queue. This is fine, and keeps system latency low. The relationship of which Subflow is getting which specific grant will also change. This effect appears as low frequency jitter that the far end must manage.

When the CM is receiving Unsolicited Grants and detects no activity on any of its Subflows, it will send one packet with the Service Flow EH Element with the UGS_Parm field set to zero grants, and then cease transmission. The CMTS will switch from UGS mode to Real Time Polling mode.

It is not necessary for the CMTS to separately monitor packet activity since the CM does this already. Worst case, if the CMTS misses the last packet which indicated zero grants, the CMTS and CM would be back in sync at the beginning of the next talk spurt. Because of this scenario, when the CM goes from inactive to active, the CM must be able to restart transmission with either Polled Requests or Unsolicited Grants.

M.2.4 Example




Figure M-2. VAD Start-Up and Stop

Figure M-2 shows an example of a single G.711 (64 kbps) voice call with a packet size of 10 ms, and a receive jitter buffer that requires a minimum of 20 ms of voice (thus 2 packets) before it will begin playback.

Assume voice begins at time zero. After a nominal processing delay and a 10 ms packetization delay, the DSP CODEC generates voice packets which are then transferred to the upstream transmit queue. The next Polled Request is used which results in the start of the Unsolicited Grants some time later. Additional Unsolicited Grants are immediately issued to clear out the upstream queue.

These packets traverse the network and arrive at the receive jitter buffer. The 20 ms minimum jitter buffer is met when the second packet arrives. Because the packets arrived close together, only an additional few milliseconds of latency has been added. After a nominal processing delay, playout begins.

When the voice spurt ends, the CM sends one remaining packet with no payload, and with the Service Flow EH Element with the UGS_Parm field set to zero grants. Some time later, UGS stops, and Real Time Polling begins.

M.2.5 Talk Spurt Grant Burst

The extra burst of Unsolicited Grants when a flow becomes active is necessary because the jitter buffer at the receiving CODEC typically waits to have a minimum amount of voice samples before beginning the playout. Any delay between the arrival of these initial packets will add to the final latency of the phone call. Thus, the sooner the CMTS recognizes that the CM has packets to send and can empty the CM's buffer, the sooner those packet will reach the receiver, and the lower the latency that will be incurred in the phone call.

It is an indeterminate problem as to how many grants must be burst. When the CM makes its request for an additional grant, one voice packet has already accumulated. The CM has no idea how many extra grants to request as it has no idea of the round trip response time it will receive from the CMTS, and thus how many packets may accumulate. The CMTS has a better idea, although it does not know the far end jitter buffer requirements.

The solution is for the CMTS to choose the burst size, and burst these grants close together at the beginning of the talk spurt. This occurs when moving from Real Time Polling to UGS, and when increasing the number of UGS Grants per Interval.

A typical start-up latency that will be introduced by the Request to Grant response time is shown in Table M-1.

Variable		Example Value	
1	The time taken from when the voice packet was created to the time that voice packet arrives in the CM upstream queue.	0 - 1	ms
2	The time until a polled request is received. The worst case time is the Polled Request Interval.	0 - 5	ms
3	The Request-Grant response time of the CMTS. This value is affected by MAP length and the number of outstanding MAPS.	5 - 15	ms
4	The round trip delay of the HFC plant including the downstream interleaving delay.	1 - 5	ms
Total		6 - 26	ms

Table M-1. Example Request to Grant Response Time

This number will vary between CMTS implementations, but a reasonable number of extra grants to expect from the example above would be:

UGS Interval	Extra Grants for New Talk Spurts
10 ms	2
20 ms	1
30 ms	0

Table M-2. Example Extra Grants for New Talk Spurts

Once again it is worth noting that the CMTS and CM cannot and do not associate individual Subflows with individual grants. That means that when current Subflows are active and a new Subflow becomes active, the new Subflow will immediately begin to use the existing pool of grants. This potentially reduces the start up latency of new talk spurts, but increases the latency of the other Subflows. When the burst of grants arrives, it is shared with all the Subflows, and restores or even reduces the original latency. This is a jitter component. The more Subflows that are active, the less impact that adding a new Subflow has.

M.2.6 Admission Considerations

Note that when configuring the CMTS admission control, the following factors must be taken into account.

VAD allows the upstream to be over provisioned. For example, an upstream that might normally handle 24 VoIP sessions might be over provisioned as high as 36 (50%) or even 48 (100%). Whenever there is over provisioning, there exists the statistical possibility that all upstream VoIP sessions may become active. At that time, the CMTS may be unable to schedule all the VoIP traffic. Additionally, the talk spurt grant bursts would be stretched out. CM implementations of VAD should recognize this possibility, and set a limit as to how many packets they will allow to accumulate on its queue.

Occasional saturation of the upstream during VAD can be eliminated by provisioning the maximum number of permitted VoIP sessions to be less than the maximum capacity of the upstream with all voice traffic (24 in the previous example). VAD would cause the channel usage to drop from 100% to around 40% for voice, allowing the remaining 60% to be used for data and maintenance traffic.

Appendix N. References

- [CableLabs1] Two-Way Cable Television System Characterization, Cable Television Laboratories, Inc., April 12, 1995.
- [CableLabs2] Digital Transmission Characterization of Cable Television Systems, Cable Television Laboratories, Inc., November, 1994.
- [DIX] Ethernet Protocol Version 2.0, Digital, Intel, Xerox, 1982.
- [DOCSIS3] Data-Over-Cable Service Interface Specifications, Cable Modem Termination System - Network Side Interface Specification, SP-CMTS-NSI-I01-960702.
- [DOCSIS4] Data-Over-Cable Service Interface Specifications, Cable Modem to Customer Premise Equipment Interface Specification, SP-CMCI-I02-980317.
- [DOCSIS5] Data-Over-Cable Service Interface Specifications, Operations Support System Interface Specification, SP-OSSI-I01-970403.
- [DOCSIS6] Data-Over-Cable Service Interface Specifications, Cable Modem Telephony Return Interface Specification, SP-CMTRI-I01-970804.
- [DOCSIS8] Data-Over-Cable Service Interface Specifications, Baseline Privacy Plus Interface Specification, SP-BPI-W02-981228 (in preparation).
- [DOCSIS9] Data-Over-Cable Service Interface Specifications, Radio Frequency Interface Specification, SP-RFI-I04-980724.
- [FCC15] Code of Federal Regulations, Title 47, Part 15, October 1998.
- [FCC76] Code of Federal Regulations, Title 47, Part 76, October 1998.
- [ID-CDMIB] Roeck, G., Cable-Device MIB, IETF IPCDN Internet Draft, draft-ietf-ipcdn-cable-device-mib-07.txt (in process)
- [IEEE802] IEEE Std 802-1990, Local and Metropolitan Area Networks: Overview and Architecture.
- [IEEE802.1Q] IEEE Draft Standard 802.1Q/D4. Draft Standard for Virtual Bridged Local Area Networks. December 20, 1996.
- [IMA] Internet Assigned Numbers Authority, Internet Multicast Addresses, <http://www.isi.edu/in-notes/iana/assignments/multicast-addresses>
- [IPS-SP-406] ANSI/SCTE Recommended "F" Port (Female, Indoor), Society of Cable Television Engineers
- [EIA-S542] EIA Standard 542 (1997), "Cable Television Channel Identification Plan", May 1997.
- [ISO8025] ISO 8025 (December 1987) -- Information processing systems - Open Systems Interconnection - Specification of the Basic Encoding Rules for Abstract Syntax Notation One (ASN.1).

- [ISO8802-2] ISO/IEC 8802-2: 1994 (IEEE Std 802.2: 1994) - Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 2: Logical link control
- [ISO8802-3] ISO/IEC 8802-3: 1996 (IEEE Std 802.3: 1996) - Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical sublayer specifications.
- [ISO/IEC10038] ISO/IEC 10038 (ANSI/IEEE Std 802.1D): 1993, Information technology - Telecommunications and information exchange between systems - Local area networks - Media access control (MAC) bridges.
- [ISO/IEC10039] ISO/IEC 10039:1991 Information technology – Open Systems Interconnection –Local area networks –Medium Access Control (MAC) service definition.
- [ISO/IEC15802-1] ISO/IEC 10039:1991 Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Common specifications – Part 1: Medium Access Control (MAC) service definition.
- [ITU-T H.222.0] ITU-T Recommendation H.222.0 (1995) | ISO/IEC 13818-1:1996, Information technology -- generic coding of moving pictures and associated audio information systems.
- [ITU-T J.83-B] Annex B to ITU-T Recommendation J.83 (4/97), Digital multi-programme systems for television sound and data services for cable distribution.
- [ITU-T X.25] ITU-T Recommendation X.25(03/93), Interface between data terminal equipment and data circuit-terminating equipment for terminals operating in the packet mode and connected to public data networks by dedicated circuit.
- [ITU-T Z.100] ITU-T Recommendation Z.100 (3/93) - CCITT Specification and description language (SDL).
- [NCTA] NCTA Recommended Practices for measurement on Cable Television Systems - National Cable Television Association, Washington DC, 2nd Edition, revised October, 1993.
- [PKTCBL-MGCP] PacketCable Specifications, Network-Based Call Signaling Protocol Specification, Pkt-SP-EC-MGCP-D02-990226.
- [RFC-791] Postel, J., Internet Protocol, IETF RFC-791 (MIL STD 1777), September, 1981.
- [RFC-826] Plummer, D., Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48-bit Ethernet address for transmission on Ethernet hardware, November, 1982.
- [RFC-868] Harrenstien, K., and Postel, J., Time Protocol, IETF RFC-868, May 1983.
- [RFC-1042] Postel, J., and Reynolds, J., A Standard for the Transmission of IP Datagrams over IEEE 802 Networks, IETF RFC-1042, February, 1988.
- [RFC-1058] Hedrick, C., Routing Information Protocol, IETF RFC-1058, June, 1988.

- [RFC-1123] Braden, R., Requirements for Internet Hosts -- Application and Support, IETF RFC-1123, October 1989.
- [RFC-1157] Schoffstall, M., Fedor, M., Davin, J. and Case, J., A Simple Network Management Protocol (SNMP), IETF RFC-1157, May, 1990.
- [RFC-1350] Sollings, K., The TFTP Protocol (Revision 2), IETF RFC-1350, July, 1992.
- [RFC-1493] Definitions of Managed Objects for Bridges. E. Decker, P. Langille, A. Rijsinghani, & K. McCloghrie. July 1993. (Obsoletes RFC1286)
- [RFC-1633] Braden, R., Clark, D., and Shenker, S., Integrated Services in the Internet Architecture: An Overview, IETF RFC-1633, June, 1994.
- [RFC-1812] Baker, F., Requirements for IP Version 4 Routers, IETF RFC-1812. June, 1995.
- [RFC-2104] Krawczyk, H., Bellare, M., and Canetti, R., HMAC: Keyed-Hashing for Message Authentication, IETF RFC-2104, February, 1997.
- [RFC-2131] Droms, R., Dynamic Host Configuration Protocol, IETF RFC-2131, March, 1997.
- [RFC-2132] Alexander, S., and Droms, R., DHCP Options and BOOTP Vendor Extensions, IETF RFC-2132, March, 1997.
- [RFC-2210] Wroclawski, J., The Use of RSVP with the IETF Integrated Services, IETF RFC-2210, September, 1997.
- [RFC-2211] Wroclawski, J., Specification of the Controlled-Load Network Element Service, IETF RFC-2211, September, 1997.
- [RFC-2212] Shenker, S., Partridge, C., and Guerin, R., Specification of Guaranteed Quality of Service, IETF RFC-2212, September, 1997.
- [RFC-2236] Fenner, W., Internet Group Management Protocol, Version 2, IETF RFC-2236, November 1997.
- [RFC-2349] Malkin, G. and Harkin, A., TFTP Timeout Interval and Transfer Size Options, IETF RFC-2349, May 1998.
- [SHA] NIST, FIPS PUB 180-1: Secure Hash Standard, April 1995.
- [SMS] The Spectrum Management Application (SMA) and the Common Spectrum Management Interface (csmi), Time Warner Cable, December 24, 1995.

This page intentionally left blank.

Appendix O. Glossary

Active Service Flow

An admitted Service Flow from the CM to the CMTS which is available for packet transmission.

Address Resolution Protocol (ARP)

A protocol of the IETF for converting network addresses to 48-bit Ethernet addresses.

Admitted Service Flow

A Service Flow, either provisioned or dynamically signaled, which is authorized and for which resources have been reserved but is not active.

American National Standards Institute (ANSI)

A US standards body.

ANSI

See American National Standards Institute.

ARP

See Address Resolution Protocol.

Asynchronous Transfer Mode (ATM)

A protocol for the transmission of a variety of digital signals using uniform 53-byte cells.

ATM

See Asynchronous Transfer Mode.

Authorization Module

The authorization module is an abstract module that the CMTS can contact to authorize Service Flows and Classifiers. The authorization module tells the CMTS whether the requesting CM is authorized for the resources it is requesting.

Availability

In cable television systems, availability is the long-term ratio of the actual RF channel operation time to scheduled RF channel operation time (expressed as a percent value) and is based on a bit error rate (BER) assumption.

Bandwidth Allocation Map

The MAC Management Message that the CMTS uses to allocate transmission opportunities to CMs.

BPDU

See Bridge Protocol Data Unit.

Bridge Protocol Data Unit (BDU)

Spanning tree protocol messages as defined in [ISO/IEC10038].

Broadcast Addresses

A predefined destination address that denotes the set of all data network service access points.

Burst Error Second

Any Errored Second containing at least 100 errors.

Cable Modem (CM)

A modulator-demodulator at subscriber locations intended for use in conveying data communications on a cable television system.

Cable Modem Termination System (CMTS)

Cable modem termination system, located at the cable television system headend or distribution hub, which provides complementary functionality to the cable modems to enable data connectivity to a wide-area network.

Cable Modem Termination System - Network Side Interface (CMTS-NSI)

The interface, defined in [DOCSIS3], between a CMTS and the equipment on its network side.

Cable Modem to CPE Interface (CMCI)

The interface, defined in [DOCSIS4], between a CM and CPE.

Carrier Hum Modulation

The peak-to-peak magnitude of the amplitude distortion relative to the RF carrier signal level due to the fundamental and low-order harmonics of the power-supply frequency.

Carrier-to-Noise Ratio (C/N or CNR)

The square of the ratio of the root mean square (rms) of the voltage of the digitally-modulated RF carrier to the rms of the continuous random noise voltage in the defined measurement bandwidth. (If not specified explicitly, the measurement bandwidth is the symbol rate of the digital modulation; for video it is 4 MHz).

Classifier

A set of criteria used for packet matching according to TCP, UDP, IP, LLC, and/or 802.1P/Q packet fields. A classifier maps each packet to a Service Flow. A Downstream Classifier is used by the CMTS to assign packets to downstream service flows. An Upstream Classifier is used by the CM to assign packets to upstream service flows.

CM

See Cable Modem.

CMCI

See Cable Modem to CPE Interface.

CMTS

See Cable Modem Termination System.

CMTS-NSI

See Cable Modem Termination System - Network Side Interface.

Composite Second Order Beat (CSO)

The peak of the average level of distortion products due to second-order non-linearities in cable system equipment.

Composite Triple Beat (CTB)

The peak of the average level of distortion components due to third-order non-linearities in cable system equipment.

CPE

See Customer Premises Equipment.

Cross-Modulation

A form of television signal distortion where modulation from one or more television channels is imposed on another channel or channels.

Customer

See End User.

Customer Premises Equipment (CPE)

Equipment at the end user's premises; MAY be provided by the end user or the service provider.

Data Link Layer

Layer 2 in the Open System Interconnection (OSI) architecture; the layer that provides services to transfer data over the transmission link between open systems.

DHCP

See Dynamic Host Configuration Protocol.

Distribution Hub

A location in a cable television network which performs the functions of a Headend for customers in its immediate area, and which receives some or all of its television program material from a Master Headend in the same metropolitan or regional area.

Downstream

In cable television, the direction of transmission from the headend to the subscriber.

Drop Cable

Coaxial cable that connects to a residence or service location from a directional coupler (tap) on the nearest coaxial feeder cable.

Dynamic Host Configuration Protocol (DHCP)

An Internet protocol used for assigning network-layer (IP) addresses.

Dynamic Range

The ratio between the greatest signal power that can be transmitted over a multichannel analog transmission system without exceeding distortion or other performance limits, and the least signal power that can be utilized without exceeding noise, error rate or other performance limits.

ECN

See Engineering Change Notice.

ECO

See Engineering Change Order.

ECR

See Engineering Change Request.

Electronic Industries Association (EIA)

A voluntary body of manufacturers which, among other activities, prepares and publishes standards.

End User

A human being, organization, or telecommunications system that accesses the network in order to communicate via the services provided by the network.

Engineering Change Notice

The final step in the procedure to change specifications.

Engineering Change Order

The second step in the procedure to change specifications. DOCSIS posts ECO to web site EC table and ECO page (with indication of ECO Comment Deadline). DOCSIS issues ECO announcement to DOCSIS-announce and working group mail lists (with indication of ECO Comment Deadline).

Engineering Change Request

The first step in the procedure to change specifications. DOCSIS issues ECR number, posts to web site EC table and ECR page. DOCSIS sends ECR to subject area working group mail list (and author).

Errored Second

Any 1-sec interval containing at least one bit error.

Extended Subsplit

A frequency division scheme that allows bidirectional traffic on a single coaxial cable. Reverse path signals come to the headend from 5 to 42 MHz. Forward path signals go from the headend from 50 or 54 MHz to the upper frequency limit.

FDDI

See Fiber Distributed Data Interface.

Feeder Cable

Coaxial cables that run along streets within the served area and connect between the individual taps which serve the customer drops.

Fiber Distributed Data Interface (FDDI)

A fiber-based LAN standard.

Fiber Node

A point of interface between a fiber trunk and the coaxial distribution.

Forward Channel

The direction of RF signal flow away from the headend toward the end user; equivalent to Downstream.

Group Delay

The difference in transmission time between the highest and lowest of several frequencies through a device, circuit or system.

Guard Time

Minimum time allocated between bursts in the upstream referenced from the symbol center of the last symbol of a burst to the symbol center of the first symbol of the following burst. The guard time should be at least the duration of five symbols plus the maximum system timing error.

Harmonic Related Carrier (HRC)

A method of spacing television channels on a cable television system in exact 6-MHz increments, with all carrier frequencies harmonically related to a common reference.

Headend

The central location on the cable network that is responsible for injecting broadcast video and other signals in the downstream direction. See also Master Headend, Distribution Hub.

Header

Protocol control information located at the beginning of a protocol data unit.

HFC

See Hybrid Fiber/Coax (HFC) System.

High Frequency (HF)

Used in this document to refer to the entire subsplit (5-30 MHz) and extended subsplit (5-42 MHz) band used in reverse channel communications over the cable television network.

High Return

A frequency division scheme that allows bi-directional traffic on a single coaxial cable. Reverse channel signals propagate to the headend above the downstream passband.

Hum Modulation

Undesired modulation of the television visual carrier by the fundamental or low-order harmonics of the power supply frequency, or other low-frequency disturbances.

Hybrid Fiber/Coax (HFC) System

A broadband bidirectional shared-media transmission system using fiber trunks between the headend and the fiber nodes, and coaxial distribution from the fiber nodes to the customer locations.

ICMP

See Internet Control Message Protocol.

IE

See Information Element.

IEEE

See Institute of Electrical and Electronic Engineers.

IETF

See Internet Engineering Task Force.

IGMP

See Internet Group Management Protocol.

Incremental Related Carriers (IRC)

A method of spacing NTSC television channels on a cable television system in which all channels except 5 and 6 correspond to the standard channel plan, used to reduce composite triple beat distortions.

Institute of Electrical and Electronic Engineers (IEEE)

A voluntary organization which, among other things, sponsors standards committees and is accredited by the American National Standards Institute.

International Electrotechnical Commission (IEC)

An international standards body.

International Organization for Standardization (ISO)

An international standards body, commonly known as the International Standards Organization.

Internet Control Message Protocol (ICMP)

An Internet network-layer protocol.

Internet Engineering Task Force (IETF)

A body responsible, among other things, for developing standards used in the Internet.

Internet Group Management Protocol (IGMP)

A network-layer protocol for managing multicast groups on the Internet

Impulse Noise

Noise characterized by non-overlapping transient disturbances.

Information Element

The fields that make up a MAP and define individual grants, deferred grants, etc.

Internet Protocol (IP)

An Internet network-layer protocol.

Interval Usage Code

A field in MAPs and UCDs to link burst profiles to grants.

IP

See Internet Protocol.

IUC

See Interval Usage Code.

Latency

The time, expressed in quantity of symbols, taken for a signal element to pass through a device.

Layer

A subdivision of the Open System Interconnection (OSI) architecture, constituted by subsystems of the same rank

LLC

See Logical Link Control (LLC) procedure.

Local Area Network (LAN)

A non-public data network in which serial transmission is used for direct data communication among data stations located on the user's premises.

Logical Link Control (LLC) procedure

In a local area network (LAN) or a Metropolitan Area Network (MAN), that part of the protocol that governs the assembling of data link layer frames and their exchange between data stations, independent of how the transmission medium is shared.

MAC

See Media Access Control (MAC) procedure.

MAC Service Access Point

See Section 6.1.2.2.

MAP

See Bandwidth Allocation Map.

Master Headend

A headend which collects television program material from various sources by satellite, microwave, fiber and other means, and distributes this material to Distribution Hubs in the same metropolitan or regional area. A Master Headend MAY also perform the functions of a Distribution Hub for customers in its own immediate area.

Mean Time to Repair (MTTR)

In cable television systems, the MTTR is the average elapsed time from the moment a loss of RF channel operation is detected up to the moment the RF channel operation is fully restored.

Media Access Control (MAC) address

The "built-in" hardware address of a device connected to a shared medium.

Media Access Control (MAC) procedure

In a subnetwork, that part of the protocol that governs access to the transmission medium independent of the physical characteristics of the medium, but taking into account the topological aspects of the subnetworks, in order to enable the exchange of data between nodes. MAC procedures include framing, error protection, and acquiring the right to use the underlying transmission medium.

Media Access Control (MAC) sublayer

The part of the data link layer that supports topology-dependent functions and uses the services of the Physical Layer to provide services to the logical link control (LLC) sublayer.

Micro-reflections

Echoes in the forward transmission path due to departures from ideal amplitude and phase characteristics.

Mid Split

A frequency division scheme that allows bi-directional traffic on a single coaxial cable. Reverse channel signals propagate to the headend from 5 to 108 MHz. Forward path signals go from the headend from 162 MHz to the upper frequency limit. The diplex crossover band is located from 108 to 162 MHz.

Mini-Slot

A "mini-slot" is an integer multiple of 6.25-microsecond increments. The relationship between mini-slots, bytes and time ticks is described in Section 7.3.4.

Moving Picture Experts Group (MPEG)

A voluntary body which develops standards for digital compressed moving pictures and associated audio.

MPEG

See Moving Picture Experts Group.

MSAP

See MAC Service Access Point.

Multipoint Access

User access in which more than one terminal equipment is supported by a single network termination.

Multipoint Connection

A connection among more than two data network terminations.

National Cable Television Association (NCTA)

A voluntary association of cable television operators which, among other things, provides guidance on measurements and objectives for cable television systems in the USA.

National Television Systems Committee (NTSC)

Committee which defined the analog color television broadcast standard used today in North America.

Network Layer

Layer 3 in the Open System Interconnection (OSI) architecture; the layer that provides services to establish a path between open systems.

Network Management

The functions related to the management of data link layer and physical layer resources and their stations across the data network supported by the hybrid fiber/coax system.

Open Systems Interconnection (OSI)

A framework of ISO standards for communication between different systems made by different vendors, in which the communications process is organized into seven different categories that are placed in a layered sequence based on their relationship to the user. Each layer uses the layer immediately below it and provides a service to the layer above. Layers 7 through 4 deal with end-to-end communication between the message source and destination, and layers 3 through 1 deal with network functions.

Organizationally Unique Identifier (OUI)

A 3-octet IEEE assigned identifier that can be used to generate Universal LAN MAC addresses and Protocol Identifiers per ANSI/IEEE Std 802 for use in Local and Metropolitan Area Network applications.

OSI

See Open Systems Interconnection.

OUI

See Organization Unique Identifier.

Packet Identifier (PID)

A unique integer value used to identify elementary streams of a program in a single- or multi-program MPEG-2 stream.

Partial Grant

A grant that is smaller than the corresponding bandwidth request from the CM.

Payload Header Suppression

The suppression of the header in a payload packet. (e.g. the suppression of the Ethernet header in forwarded packets)

Payload Unit Start Indicator (PUSI)

A flag in an MPEG header. A value of 1 indicates the presence of a pointer field as the first byte of the payload.

PHS

See Payload Header Suppression.

PHY

See Physical (PHY) Layer.

Physical (PHY) Layer

Layer 1 in the Open System Interconnection (OSI) architecture; the layer that provides services to transmit bits or groups of bits over a transmission link between open systems and which entails electrical, mechanical and handshaking procedures.

Physical Media Dependent (PMD) Sublayer

A sublayer of the Physical Layer which is concerned with transmitting bits or groups of bits over particular types of transmission link between open systems and which entails electrical, mechanical and handshaking procedures.

PID

See Packet Identifier.

PMD

See Physical Media Dependent (PMD) Sublayer.

Primary Service Flow

All CMs have a Primary Upstream Service Flow and a Primary Downstream Service Flow. They ensure that the CM is always manageable and they provide a default path for forwarded packets that are not classified to any other Service Flow

Program-Specific Information (PSI)

In MPEG-2, normative data necessary for the demultiplexing of Transport Streams and the successful regeneration of programs.

Program Stream

In MPEG-2, a multiplex of variable-length digital video and audio packets from one or more program sources having a common time-base.

Protocol

A set of rules and formats that determines the communication behavior of layer entities in the performance of the layer functions.

Provisioned Service Flow

A Service Flow that has been provisioned as part of the Registration process, but has not yet been activated or admitted. It may still require an authorization exchange with a policy module or external policy server prior to admission.

PSI

See Program-Specific Information.

QAM

See Quadrature Amplitude Modulation.

QoS Parameter Set

The set of Service Flow Encodings that describe the Quality of Service attributes of a Service Flow or a Service Class. (Refer to Appendix C.2.2.5)

QPSK

See Quadrature Phase-Shift Keying.

Quadrature Amplitude Modulation (QAM)

A method of modulating digital signals onto a radio-frequency carrier signal involving both amplitude and phase coding.

Quadrature Phase-Shift Keying (QPSK)

A method of modulating digital signals onto a radio-frequency carrier signal using four phase states to code two digital bits.

Radio Frequency (RF)

In cable television systems, this refers to electromagnetic signals in the range 5 to 1000 MHz.

Request For Comments (RFC)

A technical policy document of the IETF; these documents can be accessed on the World Wide Web at <http://ds.internic.net/ds/rfcindex.html>.

Return Loss

The parameter describing the attenuation of a guided wave signal (e.g., via a coaxial cable) returned to a source by a device or medium resulting from reflections of the signal generated by the source.

Reverse Channel

The direction of signal flow towards the headend, away from the subscriber; equivalent to Upstream.

RFC

See Request for Comments.

Routing Information Protocol (RIP)

A protocol of the IETF for exchanging routing information about IP networks and subnets.

SAID

See Security Association Identifier.

Service Access Point (SAP)

The point at which services are provided by one layer, or sublayer to the layer immediately above it.

Security Association Identifier

A Baseline Privacy security identifier between a CMTS and a CM.

Service Data Unit (SDU)

Information that is delivered as a unit between peer service access points

Service Class

A set of queuing and scheduling attributes that is named and that is configured at the CMTS. A Service Class is identified by a Service Class Name. A Service Class has an associated QoS Parameter Set.

Service Class Name

An ASCII string by which a Service Class may be referenced in modem configuration files and protocol exchanges.

Service Flow

A MAC-layer transport service which:

- Provides unidirectional transport of packets from the upper layer service entity to the RF;
- Shapes, polices, and prioritizes traffic according to QoS traffic parameters defined for the Flow.

Service Flow Identifier (SFID)

An identifier assigned to a service flow by the CMTS. [32 bits]

Service Identifier (SID)

A Service Flow Identifier assigned by the CMTS (in addition to a Service Flow Identifier) to an Active or Admitted Upstream Service Flow. [14 bits]

Service Flow Reference

A message parameter in Configuration Files and Dynamic Service MAC messages used to associate Classifiers and other objects in the message with the Service Flow Encodings of a requested Service Flow.

SID

See Service Identifier.

Simple Network Management Protocol (SNMP)

A network management protocol of the IETF.

SMS

See Spectrum Management System.

SNAP

See Subnetwork Access Protocol.

SNMP

See Simple Network Management Protocol.

Spectrum Management System (SMS)

A system, defined in [SMS], for managing the RF cable spectrum.

Sublayer

A subdivision of a layer in the Open System Interconnection (OSI) reference model.

Subnetwork

Subnetworks are physically formed by connecting adjacent nodes with transmission links.

Subnetwork Access Protocol (SNAP)

An extension of the LLC header to accommodate the use of 802-type networks as IP networks.

Subscriber

See End User.

Subsplit

A frequency-division scheme that allows bi-directional traffic on a single cable. Reverse path signals come to the headend from 5 to 30 (up to 42 on Extended Subsplit systems) MHz. Forward path signals go from the headend from 50 or 54 MHz to the upper frequency limit of the cable network.

Subsystem

An element in a hierarchical division of an Open System that interacts directly with elements in the next higher division or the next lower division of that open system.

Systems Management

Functions in the application layer related to the management of various open systems Interconnection (OSI) resources and their status across all layers of the OSI architecture.

TFTP

See Trivial File-Transfer Protocol.

Tick

6.25-microsecond time intervals that are the reference for upstream mini-slot definition and upstream transmission times.

Tilt

Maximum difference in transmission gain of a cable television system over a given bandwidth (typically the entire forward operating frequency range).

TLV

See Type/Length/Value.

Transit Delay

The time difference between the instant at which the first bit of a PDU crosses one designated boundary, and the instant at which the last bit of the same PDU crosses a second designated boundary.

Transmission Control Protocol (TCP)

A transport-layer Internet protocol which ensures successful end-to-end delivery of data packets without error.

Transmission Convergence Sublayer

A sublayer of the Physical Layer that provides an interface between the Data Link Layer and the PMD Sublayer.

Transmission Link

The physical unit of a subnetwork that provides the transmission connection between adjacent nodes.

Transmission Medium

The material on which information signals may be carried; e.g., optical fiber, coaxial cable, and twisted-wire pairs.

Transmission System

The interface and transmission medium through which peer physical layer entities transfer bits.

Transmit On/Off Ratio

In multiple-access systems, the ratio between the signal powers sent to line when transmitting and when not transmitting.

Transport Stream

In MPEG-2, a packet-based method of multiplexing one or more digital video and audio streams having one or more independent time bases into a single stream.

Trivial File-Transfer Protocol (TFTP)

An Internet protocol for transferring files without the requirement for user names and passwords that is typically used for automatic downloads of data and software.

Trunk Cable

Cables that carry the signal from the headend to groups of subscribers. The cables can be either coaxial or fiber depending on the design of the system.

Type/Length/Value (TLV)

An encoding of three fields, in which the first field indicates the type of element, the second the length of the element, and the third field the value.

Upstream

The direction from the subscriber location toward the headend.

Upstream Channel Descriptor (UCD)

The MAC Management Message used to communicate the characteristics of the upstream physical layer to the cable modems.

Appendix P. Acknowledgment

Rich Woundy of American Internet Corporation (now part of Cisco) led the Quality of Service working group that contributed additions to Sections 6 and 8 as well as Appendix C. Mike Patrick of Motorola and Burcak Beser of 3com were instrumental members of this group and contributed text and ideas for many portions of this specification.

Tom Quigley and Lisa Denney of Broadcom Corporation as well as John Ulm of Nortel Networks, Inc. Data Over Cable Division spearheaded the effort to incorporate fragmentation into the specification.

John Ulm, David Unger, Wilson Sawyer and Gerry White of Nortel Networks, Inc. Data Over Cable Division wrote the original versions of Section 3 and Section 5 through Section 10, as well as Appendices A through E, G, H, I, and K. Wilson Sawyer authored the original version of the Theory of Operations section of Chapter 8. John Ulm revised all of the state transition diagrams in Section 9.4 and authored the new Appendix L.

Steven Anderson, Tom Kolze, Victor Hou and Clive Holborow of General Instrument Corporation wrote Section 4 and Appendix F.

John Chapman of Cisco Systems wrote the Dynamic Service Change and the Payload Header Suppression sections as well as contributing the idea for Unsolicited Grants with Activity Detection — including authoring Appendix M.

John Pickens and James Yee of Com21 made significant contributions to the Theory of Operations section of Chapter 8 and the QoS Parameters in Appendix C. as well as updating Appendix E.

Mike St. Johns of @Home Corporation spearheaded the addition of IGMP support to the specification. The contributions of Paul Gray and Mike Patrick of Motorola Corporation provided an additional compatible approach.

CableLabs and the cable industry as a whole are grateful to these individuals and organizations for their contributions.

This page intentionally left blank.

Appendix Q. Revisions

Q.1 ECNs Included in SP-RFiv1.1-I02-990731

Table Q-1. Incorporated ECN Table

ECN	Date Accepted
rfl-n-99019	07/07/99
rfl-n-99025	05/19/99
rfl-n-99035	06/02/99
rfl-n-99039	06/23/99
rfl-n-99040	06/23/99
rfl-n-99043	06/30/99
rfl-n-99048	07/07/99
rfl-n-99049	07/07/99
rfl-n-99050	07/07/99
rfl-n-99051	07/07/99
rfl-n-99052	07/07/99
rfl-n-99053	07/07/99
rfl-n-99054	07/07/99
rfl-n-99056	07/07/99

This page intentionally left blank.

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ **BLACK BORDERS**
- ☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- ☐ **FADED TEXT OR DRAWING**
- ☐ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- ☐ **SKEWED/SLANTED IMAGES**
- ☐ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- ☐ **GRAY SCALE DOCUMENTS**
- ☐ **LINES OR MARKS ON ORIGINAL DOCUMENT**
- ☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- ☐ **OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.